// Amreesh Phokeer (AFRINIC)
// Kevin Chege (Internet Society)
// Josiah Chavula (University of Cape Town)
// Ahmed Elmokashfi (Simula Research Lab)
// Assane Gueye (CMU-Africa)

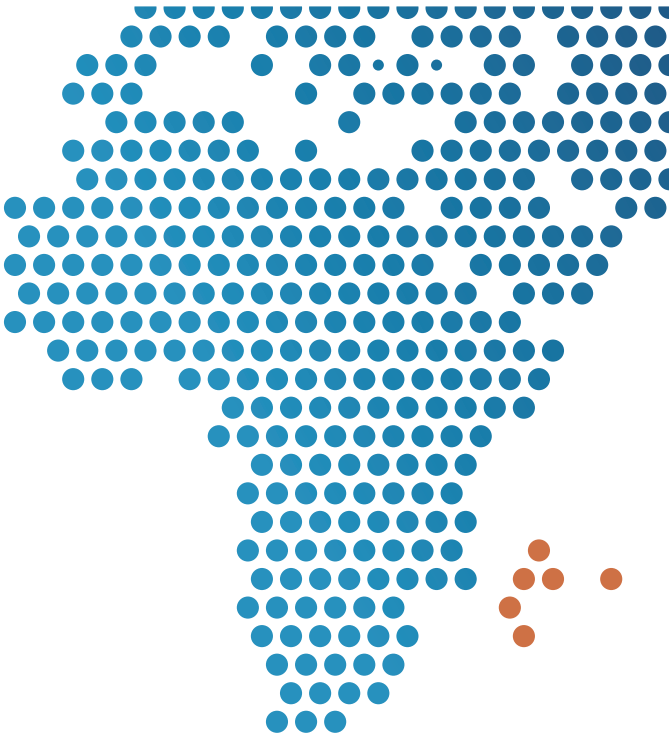

April 2021 (v1.0)

Measuring Internet Resilience in Africa (MIRA)

Project Overview

Internet Society

AFRINIC
The Internet Numbers Registry for Africa

Abstract

The Internet plays a critical role in society today. The COVID-19 pandemic has further underlined the importance of reliable Internet connectivity for everyone. Unfortunately, not all countries have achieved sufficient maturity in terms of having reliable Internet infrastructure. In particular, low-income countries usually have under-provisioned networks and lack proper cable infrastructure or redundant interconnection systems. In these countries (or regions), significant Internet outages occur when there is a cable break or power failure. This impacts the whole Internet ecosystem, which can result in significant revenue loss for the digital economy. Additionally, many low-income countries do not have the capacity to thoroughly audit their Internet infrastructure and, in many cases, they have not developed or adopted best practices for building resilient Internet infrastructure. The Measuring Internet Resilience in Africa (MIRA) project is a joint initiative between the African Network Information Centre (AFRINIC) and the Internet Society, with the goal of evaluating the capability of a country to provide a stable and reliable means of Internet connectivity at all times. Based on the results, we will provide recommendations in the form of best practices that could help networks or countries achieve higher Internet resilience.

1. Introduction

One of the 13 principles of the African Declaration of Internet rights and freedom[1] is the *Security, Stability and Resilience of the Internet*. This principle implies that everyone has the right to enjoy secure and reliable connectivity to the Internet regardless of the size and location of their network. The COVID-19 pandemic has shown the importance of the Internet and Internet-enabled services to society and has further demonstrated how critical it is to build networks that are resilient. However, many African networks are frequently subject to many forms of disruption, such as power failures, cable breaks, (un)intentional shutdowns, and other security incidents [1]. In some instances, the outages are caused by accident, either due to poor engineering or lack of redundant infrastructure. In some cases, the disruptions are due to state-sponsored shutdowns, especially during election periods [2]. In some cases, shutdowns occur because operators are not held accountable and so are not incentivized to invest in their infrastructure to make it resilient. Whether intended or not, Internet disruptions can have a considerable impact on society and the economy [3].

Recent studies have highlighted the diversity in the quality of Internet connectivity between and within African countries [4]–[6]

However, many of Africa's Internet challenges remain uninvestigated. Assertions are mostly based on unwritten anecdotal knowledge informally shared across operational communities, such as at the Africa Internet Summit (AIS) and the Africa Peering and Interconnection Forum (AfPIF). This deficiency makes it difficult to formulate evidence-based solutions or evaluate the success of newly deployed interconnections and investment schemes. A survey carried out by AFRINIC in 2019 [7] showed that Internet measurement is not a common practice in Africa. Lack of sufficient measurements in African countries makes it very challenging to accurately determine the problem areas that need to be addressed in order to improve the reliability and resiliency of the Internet in Africa.

The MIRA project will establish a framework called the "Internet Resilience framework" that will evaluate the ability of a network (and by aggregate, a country) to provide stable and reliable means of connectivity

[1] https://africaninternetrights.org

to the Internet. The framework will be based on analysis of empirical data (both primary and secondary) collected from networks and countries in Africa. Based on these results, we plan to identify and outline the best practices required for creating a more resilient national and regional interconnection system that, if implemented by Internet Service Providers (ISPs) and network operators, could strengthen and safeguard the Internet infrastructure from disruption.

1.1 // Context

This project is a joint initiative from African Network Information Centre (AFRINIC) and the Internet Society. This project falls under the AFRINIC AIM (Africa Internet Measurement) program[2] and within the Internet Society's "Measuring the Internet"[3] Project. AFRINIC and the Internet Society will be collaborating with other researchers as described in Section 6.4 to Measure Internet Resilience in Africa (MIRA). In summary we want to:

1. Collect and analyze empirical data to determine current levels of Internet resilience in African countries.
2. Further develop Internet measurement infrastructure in Africa by increasing the number of measurement vantage points that are active in Africa.
3. Present the data for users at all levels (policy makers, engineers, network operators, decision makers, Internet users, etc.).

1.2 // Target Audience

The outcome of this project will be used to inform decision-makers in two main categories:

- Network operators and Internet Service Providers (ISPs) seeking to improve the resilience of their infrastructure
- National Regulatory Authorities (NRAs) defining the legal and operational environments of Internet ecosystems in their respective countries.

This project and its findings will have wider applicability and may be of interest to a larger audience, including consumer and industry groups, academic and industrial research labs, as well as standardization bodies.

1.3 // Definition And Scope

Throughout this project, we aim to investigate the threats and obstacles (both internal and external) that impact Internet infrastructure, and the mechanisms for increasing the overall resilience of Internet services. That is, *the ability of a network to maintain an acceptable level of service in the event of an outage or during crises* [8]–[13]. The same principle applies to a country's Internet resilience and this project will evaluate each country's ability to continue providing a best-effort service during a crisis.

As shown in Figure 1, "Internet resilience" encompasses many underlying components[4], ranging from the resilience of physical Internet infrastructure and the power infrastructure to market resilience and quality of service (QoS) i.e. performance, uptime, available bandwidth, etc.

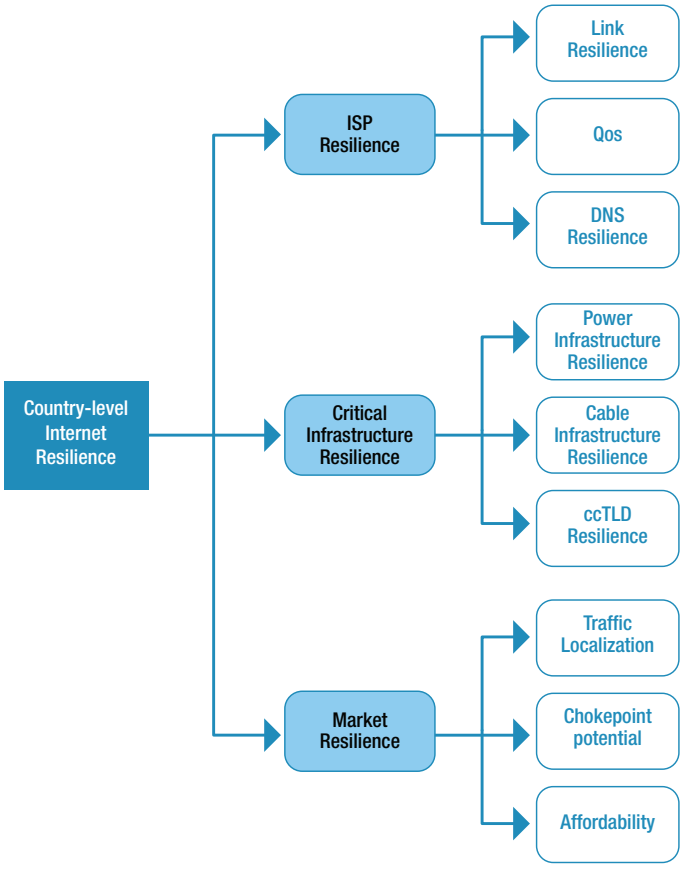

Figure 1: *Taxonomy of resilience components*

[2] https://afrinic.net/research/programmes/aim
[3] https://www.internetsociety.org/issues/measurement
[4] The diagram shows the initial scope of the project and is not a comprehensive list of all possible aspects of resilience.

Definitions

- Country-level Internet Resilience: the ability of a country's national Internet ecosystem (ISPs, regulations, physical infrastructure, market structure) to provide Internet services to its citizens at an acceptable level of service in the face of faults and challenges to normal operations.
- Critical Infrastructure Resilience: the resilience of the power infrastructure, the Internet cable infrastructure (both terrestrial and undersea), as well as the country-code Top-Level Domain (ccTLD) infrastructure.
- Market Resilience: the availability and efficiency of Internet Exchange Points (IXP) and the ability to keep local traffic local, the ability of the market to self-regulate and provide affordable prices to end-users by maintaining a diverse and competitive market.
- Network/ISP Resilience: the ability of a network to continue providing an acceptable level of service in the event of an outage or during a crisis. This resilience component is made up of various components such as the resilience of physical links, logical/peering links, performance/QoS, and DNS.

Scope and Objectives

In our evaluation framework, we will consider the following aspects:

1. The availability and stability of the physical infrastructure, which includes power stations, undersea or terrestrial fiber, landing stations, and last mile access networks.
2. The quality of service (QoS) of the network from the end user's perspective and the stability of the network in terms of reachability, throughput and latency to selected target servers.
3. The availability and performance of the Domain Name System (DNS) ecosystem.
4. The availability and efficiency of the local peering fabric as well as the ability of the country to keep local traffic local.
5. The resilience of the ISP market i.e., the level of concentration towards specific Autonomous Systems (AS) and affordability.

2. Approach and Objectives

In this project, we will collect inter-connectivity data for Autonomous System Numbers (ASNs) in Africa and use the data to analyze metrics of network resilience, reliability, and performance for inter-ASN communication. Our primary definition of resilience will be based on a network's ability to tolerate, remediate, and recover from network incidents, such as those caused by device failures and fiber cuts [10], [14]. In considering the physical topology, we will study the geographic resilience from both end-to-end and ISP (ASN) perspectives. In addition, two other related metrics will be analyzed - performance and reliability. All the metrics considered in this project will be evaluated at ASN level as well as at city and country levels. From a geographic point of view, we will group routers within different levels of a vicinity such as city, country and region. This analysis will be conducted independently from ASN level topology.

2.1 // Overview of the design

As shown in Figure 2, the MIRA framework is made up of several building blocks. *The Internet Society Pulse Dashboard* will allow end-users to visualize the data collected under the MIRA project. Users will be able to build customized dashboards as per their needs. For example, users will be able to pick and choose their indices and build their own final index. The Internet Society Pulse dashboard will pull data using the MIRA API, which can also be used by programmers to retrieve data from MIRA without using the GUI. The *Analytics pipeline* is responsible for generating the indices based on data collected (primary and/or secondary).

The *Internet Resilience framework* (4.2 WP2: Internet Resilience Framework) will provide the specifications on how to process the data. The data will come through the *Data Pipeline*, which will store and aggregate both primary and secondary datasets. Primary data will be collected through fixed and mobile measurement nodes based on Raspberry Pis, which are running measurement tools - namely RIPE Atlas

probes[5] and M-Lab's Network Diagnostic Tool (NDT) client[6]. The network measurement scripts are based on existing standard tools and will be installed and orchestrated by the MIRA measurement manager. Secondary data will come from multiple open Internet data sources such as IXP routing data, RIR allocation information and Border Gateway Protocol (BGP) routing tables.

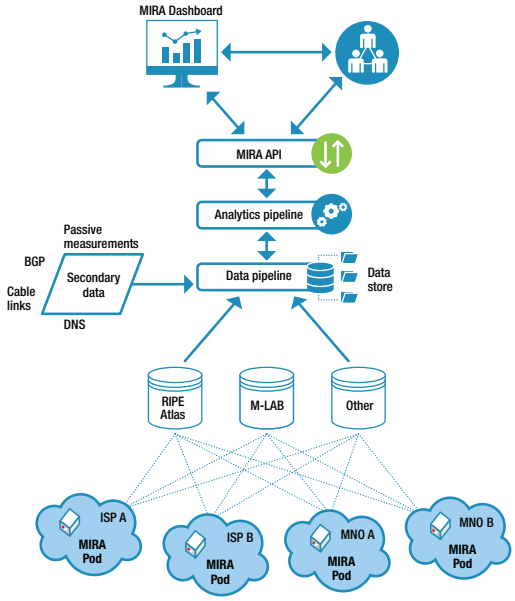

Figure 2: *MIRA Data collection and visualization workflow*

2.2 // Metrics

2.2.1 Critical Infrastructure + Path Diversity

An important approach to ensuring a resilient Internet is through increasing the diversity of network paths between any given pair of Internet hosts [10], [15]. In this context, diversity is the degree to which alternate paths share the same nodes and links [15]. To be more resilient against failures, there needs to be multiple link- and node-disjoint paths between networks [16]. Crucially, path diversity between hosts connected to different networks is determined by both the physical infrastructure (physical topology), as well as the routing policies (intra- and inter-domain) and peering policies. Inherently, path diversity is enhanced with multihoming by ASes and hosts. Thus, it is important to study both physical and logical path diversity across multiple networks.

2.2.2 Performance/QoS

Network operators generally measure network performance in terms of standard quality of service (QoS) metrics, such as throughput, delay, jitter, and packet loss. Operators may, for instance, be interested in monitoring congestion and packet loss on links within or between networks. For Internet users, the QoS metrics are useful only in terms of how they impact end-to-end communication, and how their Quality of Experience (QoE) is affected. In this case, QoE describes a user's subjective assessment of their experience when using a particular network service [20]. For example, in relation to congestion and packet loss, a user would generally only be interested in effective throughput for their applications, including determining the capability and QoE when using rich media over a given network service. If performance of a network is severely degraded, such as through congestion or device failure, it may lead to disruption of packet forwarding and cause gaps where network service is effectively not available. And this is linked to another aspect of this project—network reliability.

2.2.3 Network reliability

Network reliability is a notion that encompasses various stability metrics that are important for sustained availability and usability of network services. A key metric of network reliability is uptime, which is a measure of the percentage of time that a network service is available. The level of network uptime determines whether a user is able to access Internet services all the time. Thus, from a user's point of view, we can measure Internet reliability in terms of uptime and reachability—being able to reach any network and Internet services at any given time [17].

[5] https://atlas.ripe.net/probes
[6] https://pypi.org/project/murakami

2.3 // Measuring Path Diversity

To measure and evaluate Internet resilience with respect to path diversity, we will implement a measurement framework for identifying geographically equivalent hops in traceroute data, and then compute geo-diversity metrics for pairs of endpoints.

2.3.1 Paris-Traceroute

Through traceroute measurements, we will collect IP-level routes between selected endpoints, both within the specific geographic areas (internal diversity) as well as between different areas (external diversity). An important aspect of measuring route diversity is to identify as many alternate routes as possible between any pair of networks or endpoints. To achieve this, our measurement framework will use Paris-Traceroute[7] and will repeatedly run the measurements between each pair of endpoints.

2.3.2 Geo-clustering

Firstly, we will use multiple geolocation databases, such as RIPE IPmap[8], MaxMind[9] and IPInfo[10], as well as active measurements to determine the geographical locations (city and country level) of traceroute IP hops. We will thereafter be able to group hops based on geographical zones at multiple levels of granularity, and this will allow us to identify geographically equivalent routes. We will then determine geographical equivalence of routes and compute the geo-diversity indices between selected locations and endpoints [18].

2.3.3 Mapping Physical and Logical Topology

For insight into Internet resilience with respect to route geo-diversity, we will study the problem of mapping the logical topology onto the physical network infrastructure [22]. In this context, the logical topology comprises sets of IP or ASN nodes, with edges denoting the relationship between networks, such as customer-provider or peering relationships. On the other hand, a physical network denotes a set of physical nodes (e.g. border routers and IXPs), with edges that represent physical communication links, mostly cable systems. The goal is to evaluate and compare the diversity of physical and logical routes between networks.

2.4 // Measuring Network Reliability and Performance

We will run long-term, end-to-end active measurements in order to assess the reliability (uptime and reachability) as well as the performance of networks in Africa. Active measurements consist of sending probe packets from a source (vantage point) to a destination. These active measurements will enable us to analyze network queuing, losses, delays, throughput, routing behaviors and propagation delays. To achieve reliable and complete results, the measurements will need to be conducted between a large number of geographically distributed vantage points.

There are already multiple Internet measurement platforms with probes that can be used to repeatedly run network measurements. Some of the notable ones include SamKnows[11], Speedchecker[12], Archipelago[13], RIPE Atlas[14] and M-Lab[15]. RIPE Atlas has approximately 12k hardware probes around the globe. However, as of 2018, there were only 229 active RIPE Atlas probes in Africa. As of February of 2021, only 194 of these are connected and active. Most of these probes are deployed by network operators in their internal networks and a small number of probes are hosted in people's homes.

Speedchecker is an active measurement platform with a relatively higher number of vantage points in Africa. As of 2018, Speedchecker had up to 850 probes in Africa covering 52 countries. Just like RIPE Atlas, Speedchecker supports a wide range of network tests, including Ping (TCP/ICMP), DNS, Traceroute, and HTTP. CAIDA's Archipelago currently has 10 active monitors in Africa, and these monitors act as dedicated probes that repeatedly run network measurement tests aimed at discovering Internet topology and measuring network performance.

M-Lab is another platform that allows web users to run throughput tests from their browsers. However, M-Lab has only seven live servers in Africa that act as targets of throughput measurements, and this limits reliability of results from many vantage points. Fortunately, M-Lab maintains a vast repository of

[7] Paris traceroute is another version of a well-known network diagnosis tool. It addresses problems caused by load balancers with the initial traceroute implementation. More information here: https://paris-traceroute.net
[8] https://ipmap.ripe.net
[9] https://www.maxmind.com
[10] https://ipinfo.io
[11] https://www.samknows.com
[12] http://speedchecker.com
[13] https://www.caida.org/projects/ark
[14] https://atlas.ripe.net
[15] https://www.measurementlab.net

existing tools, and some of these can be customized as clients and servers for different measurement campaigns. For example, the Network Diagnostic Tool[16] (NDT) provides detailed packet level information along with kernel-level statistics on how a TCP connection performs on a given path. NDT can thus be used to determine the causes of slow speeds, as well as checks for proxies, NAT devices between the machine running the tests and the M-Lab server. For this project, a custom NDT client will be developed and deployed using Raspberry Pis (MIRA Pods) and will run tests against our custom servers.

Initially the MIRA project will prioritize working with M-Lab and RIPE Atlas software probes.

2.5 // Metrics Aggregation

Measuring specific network features (QoS, latency, jitter, QoE, packet loss, throughput, path diversity, etc.) has been a daily routine and a widely accepted approach to assess network performance and/or resilience. This information is supposed to help engineers, regulators, management (and other decision makers), as well as end users to take proactive actions to adjust and improve the network performance/resilience. However, to efficiently quantify the effect of our actions, we need to aggregate these individual metrics. Engineers need aggregate measurements to focus their scarce resources in improving the networks. Decision makers need these aggregate measures in order to direct their funding decisions and to establish regulation. Users need these compound metrics in their quest for a better user experience. Unfortunately, it is not trivial to aggregate multiple metrics into meaningful results. Currently, there is no systematic way to assess the performance/resilience of a network given a vector of network measurements.

The main challenge that makes it almost impossible to soundly aggregate network metrics is the lack of "ground truth". Normally, ground truth refers to information collected on location. For instance, in physics, the ground truth comes from the physical world. In computer networks, it is difficult to obtain ground truth for measures such as packet loss or throughput, etc. However, answers must be obtained in order to make informed decisions and to evaluate whether or not our actions have improved or declined network resilience.

For this project we propose to define (not discover) a ground truth for such measurements. Our approach will build on experts' opinions and, based on the latter, we shall define the aggregation methods by assigning the appropriate weightage or coefficient to the metrics collected.

2.6 // Data Sources

In this project, we will use both primary and secondary sources of data. Primary sources will come from active and passive measurement campaigns, while secondary data will be extracted from third party information, such as the BGP routing table, IXP datasets, ccTLD information, etc. For our framework to be sustainable, we need to make sure our sources of data are open, reliable and up-to-date and will keep providing data over a long period of time. This will allow us to extract trend information about a particular topic. Appendix 1 is a list of data sources, their purpose and category.

3. Program of Work

The program of work is separated into work packages (WP) namely: Stage 1, dominated by the implementation and testing of the testbed and the development of the Internet resilience framework (WP1, WP2). Stage 2 consists of collecting and analyzing network measurements and other available secondary sources of data (WP3, WP4). Stage 3 is about curating all the data captured, aggregating it, running statistical analysis and providing insights to end-users (WP7). Figure 3 shows how each WP feeds into one another.

Note: *WP5 (Mobile Measurements) and WP6 (DNS Resilience) will be undertaken at a later stage (TBD).*

3.1 // WP1: Data sources and Measurement Infrastructure

This WP will design and implement small devices called MIRA Pods. The Pods are small devices that will be dedicated to carrying out the measurements. At this point, the pods are Raspberry Pis. This will be supported by AFRINIC and the Internet Society through the purchase of the measuring devices and management of the measurement dashboard on *Internet Society Pulse*.

Note: *this will be linked into WP5: Mobile broadband measurements*

[16] https://software.internet2.edu/ndt

3.1.1 Task 1: Measurement Pods.

The first task is to design and build appropriate infrastructure to launch the measurements (e.g., for throughput or latency measurements). For this, we will leverage existing measurement systems, including RIPE Atlas or the M-Lab NDT client. These measurement devices will be called *MIRA Pods*. They should be lightweight and easy to connect into any network. The MIRA Pods will run lightweight virtualized containers (e.g., Docker or LXD) that will run the different *Measurement Tasks*. It should be fairly easy to move or clone Measurement Tasks from one pod to another. AFRINIC and the Internet Society will support deployment of the pods in a few African countries with the support of local Internet Society Chapters and other willing participants.

3.1.2 Task 2: Pods orchestration

One important aspect is the orchestration and the management of the MIRA Pods. It should be easy to spin off and schedule measurements from a centralized system. The orchestrator should have access to the MIRA Pods and retrieve the health status of the pods as well as the status of the measurements. The orchestrator also will coordinate measurement and data collection campaigns across diverse sets of Pods. It will be responsible for gathering data from Pods. Data access controls will be critical for this, allowing multiple abstractions of data to be exposed to different parties.

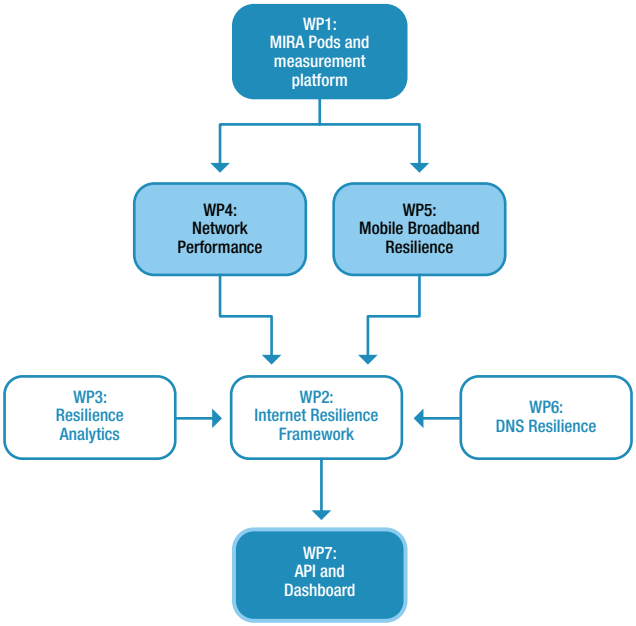

Figure 3: Information workflow between the Work Packages (WP) identified.

3.1.3 Task 3: Data pipeline and storage

The system should be able to handle primary data from the Pods and third-party data that will be used to complement the measurements. The MIRA Pods will be generating a large amount of measurement data over time. This means that the "data pipeline and storage" layer needs to: (1) handle large amount of data, (2) perform aggregation and (3), discard unnecessary data. An API should allow easy access to the data collected. In this task, we shall make use of state-of-the-art storage techniques (e.g., noSQL, Hadoop clusters, etc.) to increase scalability and ease of use.

3.2 // WP2: Internet Resilience Framework

In this WP, we will develop tools for aggregating network resilience metrics. Such aggregation will enable us to derive summary values that can quickly and intuitively give indications of network resilience. We will first establish the theoretical foundation of the aggregation by leveraging a ground truth which we plan to define.

In the first phase, we will build a simple framework aggregating the different indices using a simple formula. At a later stage, we will build a more complex framework that will harness expert opinions and then use Machine Learning to further refine the framework.

The ground truth will also serve to validate and continually recalibrate our aggregation tool. Once our metric aggregation methods are established, we will pass them onto the analytics pipeline, which defines how the different metrics (measured from different vantage points) will be fed to the aggregation modules. Finally, with sound aggregation of the resilience metrics, network operators will be able to set reference operating point and continually steer the network toward such desired reference value.

3.2.1 Task 1: Theoretical framework

In this task, we aim to develop the theoretical foundation for aggregating network resilience metrics. This shall start by identifying the metrics that can be used to effectively assess systems' overall level of resilience. Then, these metrics will be aggregated to provide a succinct summary of network resilience in the form on an *index*, called the *Internet Resilience Index*. We will build the aggregation tools by answering two main questions: (1) how to aggregate measures of similar metrics taken from vantage points within the network and measures gathered for different metrics of the same network? And (2) how

to aggregate metrics with a *zoom-in/zoom-out* effect at different levels such as city, country, and region? The aggregated metric will be mapped to numerical scores that will then be translated into a qualitative representation (such as low, medium, high, and critical) to help organizations properly assess and prioritize their decision processes.

To help build the theoretical foundation and validate our theories we will define a ground truth for network resilience metrics.

3.2.2 Task 2: Bench-marking/Ground Truth

We propose to develop and document a procedure that follows sound scientific principles and that enables the harnessing of human expert knowledge to build a resilience matrix from the chosen resilience metrics. Furthermore, it will enable the collaboration of multiple human experts such that the ground truth used for the aggregation process is based on the joint knowledge of the experts.

The Internet Society and AFRINIC maintain databases of experts in this domain. We propose to provide them our process and associated tools to enable them to assess the accuracy of our aggregation scores (to compare what their expert opinion says a score should be against what the score actually is). As human experience grows and our understanding of network resilience changes, the 'true' scores themselves change (as based on the human expert opinion). The result is that our ground truth and our metrics will change over time. By using this new knowledge, we will systematically re-calibrate and improve our aggregate methods.

3.2.3 Task 3: Analytics pipeline

The Analytics pipeline will host discrete units of computation called Analytics Modules. These modules will be defined by the Internet Resilience Framework. These will receive data streams from MIRA Pods and other sources, and will then transform them into a processed form (e.g., computing centrality from topology maps). This task will develop the APIs and framework to run these modules in a scalable manner.

Having metric aggregation methods will allow network operators and service providers to define reference points and continually steer the system toward the desired point. The resilience target reflects the requirements of end users, regulators, and other stakeholders of an acceptable operating point. The continuous steering of the network could be done by following guidance frameworks such as the NIST Network Resilience Framework [23] and Sterbenz et al. [18] resilience Framework. One such example (following NIST five step framework) could be:

1. **Protect:** by taking "proactive" measures to maintain the network into the desired operating point.
2. **Detect:** indications of network resilience degradation, which should be reflected by a decrease in score.
3. **Identify:** the metrics/parameters that are most relevant to the resilience degradation.
4. **Respond:** remediate by taking "reactive" measures to bring the system back the desired operating point.
5. **Recover:** despite continuing monitoring and protection, system might eventually suffer from disastrous failure (due to intentional or unintentional causes). This will lead to a substantial degradation in the resilience score. Appropriate measures shall be taken to recover from such disaster.

3.3 // WP3: Resilience Analytics

This work package focuses on quantifying the resilience of the African Internet at three levels.

3.3.1 Task 1: Physical topology analysis

In this task, we will assess the diversity of undersea cables that land on the continent, as well as terrestrial cables that run through the continent. We will also quantify physical elements, such as landing stations and IXPs. This work package will provide (to the Resilience Framework—WP2) metrics on diversity of cables and landing stations in respective countries and cities. This data will allow us to evaluate and compare, in quantifiable metrics, the diversity of physical routes within and between countries, cities and networks.

3.3.2 Task 2: Logical topology analysis

In this task, we will infer dependencies with respect to routing between ISPs across the continent, and their reliance on out-of-continent providers. We will rely on Paris-Traceroute measurements to collect logical routes between selected endpoints with the aim of identifying alternate routes. In addition to Traceroute data, we will also use BGP datasets to infer logical path diversity between networks and countries.

In particular, we will interface with ARDA[17], a system that synthesizes publicly available peering and routing information collected by route-collectors in Africa, notably from RouteViews[18] and Packet Clearing House[19] (PCH) Route Collectors. The key output of the work package will be index values that represent logical path diversity of networks and countries.

3.3.3 Task 3: Mapping the physical and logical topology

This task will involve mapping the logical topology onto the physical network infrastructure. The aim is to identity the physical infrastructure (cables, landing points, etc.) that are used by various ASNs. The goal is to evaluate and compare the diversity of physical and logical routes between networks. Besides identifying cables, we will also flag interconnection facilities and IXPs that are central to the stability of the Internet in the continent. Additionally, traffic flows from African countries will be associated with the respective cables. A key index out of this task will be the international physical path diversity for countries and networks.

3.3.4 Task 4: Chokepoint potential/AS hegemony

The Internet is composed of several networks that rely on each other to provide global connectivity. This means that the reachability of a network depends on the connectivity with other networks. This interdependence usually reflects the political and economic constraints within national boundaries. When there is an accrued dependence on some specific networks for global connectivity, this concentration may represent a "chokepoint". This task is about measuring the Chokepoint Potential [19] or the AS Hegemony [20] of a country.

3.4 // WP4: Network performance

3.4.1 Task 1: Access performance

One of the greatest challenges in deploying new services in Africa is the low quality of broadband provisioning in certain regions. Quantifying this is of critical importance for regulators, as well as network operators (including new entrants) who wish to best target their efforts. MIRA will provide the facility to monitor user Quality of Experience (QoE) for home broadband, public Wi-Fi and mobile providers. This will go beyond basic throughput testing and focus on capturing end user experiences across a diverse set of services. The key metric from this task will be indices representing the QoE in networks and countries.

3.4.2 Task 2: Infrastructure peering

Making decisions regarding the deployment and interconnection of infrastructure (e.g., networks, content servers) can be difficult, particularly in the developing and highly dynamic environments found throughout Africa. It has been extensively documented that peering has a positive effect on most Internet performance metrics. It is also well known that Africa lags well behind Europe in terms of its peering infrastructure. This task will use the MIRA data on the Internet Society Pulse dashboard to extract the relevant metrics to offer advice for network and content providers, who wish to find out if they should peer, who they should peer with, and where this should occur. This will have a particular focus on IXPs being deployed in the region, which is a major strategic goal of the African Union. Network operators and content providers will be able to input statistical information about their traffic and needs, and this information will then be fused with the data collected by MIRA (e.g. known paths, PoP locations, availability of IXPs), to recommend networks and locations for peering. Key metrics from this task will be indices representing the peering potency of networks. The ARDA system will be a source of information for this task.

3.4.3 Task 3: Keeping local traffic local

Keeping Internet services and content as close as possible to the end-users contributes to making the Internet a more secure and robust ecosystem. IXPs play an important role by establishing peering (traffic exchange) relationships between ISPs, content providers and Content Delivery Networks (CDN) operators, therefore allowing them to exchange traffic locally. This task will analyze the extent to which popular local content is hosted and distributed within a country and what the relative impact on QoE from the end-users' perspective is. This will be achieved by running measurements from the MIRA Pods.

3.4.4 Task 4: Network and Web Interference

A growing concern is that of web privacy, interception (e.g., HTTP header injection) and rate-limiting. This task will make use of existing tools such as OONI[20] that relies on active measurements to detect web traffic interference, as well as the technologies that underpin traffic interference. This will be done via active TCP, DNS, HTTP and TLS measurements. In this task, we will also leverage third party data from other censorship monitoring services to extract trends.

[17] African Route-collector Data Analyzer[5]
[18] http://www.routeviews.org
[19] https://www.pch.net
[20] https://ooni.org

3.5 // WP5: Mobile broadband network resilience

In this work package we seek to develop a comprehensive understanding of the usability of mobile Internet in the region, from a wide variety of vantage points.

3.5.1 Task 1: Adapt existing mobile measurement tools

This task builds on previous efforts to measure and catalogue problems with mobile Internet usage in some EU countries. In this task, we will (a) adapt an existing mobile measurement app to the African context, including integrating a list of critical information services in Africa, (b) study the usability and quality of Internet connectivity for the region (c) work with local research collaborators to conduct field measurements in their own countries.

3.5.2 Task 2: Usability of mobile broadband

In this task, we will develop a comprehensive understanding of the usability of mobile connectivity in Africa, from a wide variety of vantage points. This will provide, for the first time, a comprehensive picture on the state of Internet connectivity across the region.

3.5.3 Task 3: Informed list of critical services

In this task, we will run QoE measurements on an informed list of different websites that provide "critical information services" across the region. Such a list will provide clear taxonomy and categorization from the myriad of public and commercial information sources in the region. We will measure and catalogue the problems that will arise in accessing critical services in the region, providing pointers to policymakers on where and how to focus to minimize the digital divide in the region.

3.6 // WP6: DNS Infrastructure Resilience

Using several data sources, this WP will do a deep dive into the African DNS ecosystem. To this end, we will examine the reliability of the DNS infrastructure in the continent both for resolving out-of-continent and in-continent content. This involves examining the hosting, reliability and performance of African ccTLDs and all global DNS services with presence in Africa. Two recent studies by AFRINIC highlighted a number of issues that can affect the resilience of the DNS ecosystem.[21] [22]

3.6.1 Task 1: ccTLD resilience

Many African ccTLDs do not meet the BCP-16 recommendations by placing nameservers at both topologically and geographically diverse locations, to minimize the likelihood of a single failure disabling all of them. We will investigate whether African ccTLDs meet the BCP-16[23] minimum requirement of having at least two IPs to serve their zones.

3.6.2 Task 2: DNSSEC adoption and usage

We want to know which of the African ccTLDs have adopted DNSSEC and signed their zone. Additionally, based on the data obtained from APNIC[24], we can observe who is performing DNSSEC validation.

3.6.3 Task 3: Do53/DoH/DoT performance

We will compare the performance of Do53 (traditional DNS), DNS over TLS (DoT) and DNS over HTTPS (DoH) under different network conditions (mobile and fixed). We will uncover the causes of latency and circuitous DNS resolution paths, which amplify the performance impact of secure DNS protocols on DNS resolution time and page load time.

3.7 // WP7: MIRA Dashboard

This WP deals mainly with the dissemination of the data (after processing) through a visualization dashboard hosted on the Internet Society Pulse platform and an API to allow easy retrieval of the data.

3.7.1 Task 1: API for external usage

An API will provide easy access to the different indices collected as well as the data collected to calculate the index. Data should be made available over a reasonable amount of time to allow longitudinal analysis.

3.7.2 Task 2: Visualization dashboard

MIRA will present the data using the Internet Society Pulse platform which will be customizable based on the needs of the end users (e.g. regulators, ISPs, etc.). First the user will see a heat map of Africa with the

[21] https://afrinic.net/research/african-cctlds-technical-environment
[22] https://afrinic.net/research/african-dns-authoritative-nameservers
[23] https://tools.ietf.org/html/bcp16
[24] https://stats.labs.apnic.net/dnssec

countries colored based on the “Internet Resilience Index” that will be calculated in WP2. It will be possible to compare two or more countries side by side. The user can then dive into more granular indices that together contributed to the macro Internet Resilience Index. Participants who host measurement probes are likely to have access to more backend data than normal users.

4. Dissemination

Dissemination of the project’s results will be done via a Internet Society’s Pulse platform portal that will be openly available to all Internet users. Participants in the measurement process who host measurement probes or infrastructure will be able to obtain more technical data from the MIRA project

4.1 // Recruitment of MIRA Pod hosts

The sustainability and accuracy of the MIRA project will depend on the number of MIRA Pods carrying active measurements in Africa. Therefore, it is important to recruit and maintain a substantial pool of Pod hosts.

AFRINIC and the Internet Society, through their relationships with researchers, technical communities and the Internet Society Chapters, maintain a list of volunteer hosts. As the project evolves, we shall continue to recruit new hosts.

4.2 // Engagement and Capacity Building

Through the AFRINIC Measurement Working Group, we intend to organize a series of workshops around the broad topic of Internet measurements. We intend to cover subject areas such as network performance measurement, QoE monitoring, and Internet censorship. We will invite measurement infrastructure operators such as M-Lab, OONI, RIPE Atlas to contribute to our workshop sessions. We intend to organize the following workshops:

1. **Workshop 1:** AIS 2021, June 2021
2. **Workshop 2:** AfPIF 2021, August 2021
3. **Workshop 3:** SAFNOG 2021, November 2021

Additionally, each work package will be split into multiple scientific studies. The results will be disseminated in the form of scientific publications (conference papers or journal publications), blog posts and technical reports. Overall, we intend to make our findings very accessible to different audiences (technical and less technical).

5. Execution and Roadmap

Appendix 2 provides an overview of the tasks and when will they be executed.

5.1 // Measurement Infrastructure

The MIRA project will carry out measurements initially using M-Lab’s Murakami tool and will add in the RIPE NCC’s RIPE Atlas in the coming months. We intend to add additional tools based on future needs. Both of these tools support software clients that can be installed on a variety of operating systems and can carry out various Internet measurements. The software will be installed on small Raspberry PIs that we call MIRA pods. We chose this technology to allow uninterrupted and dedicated measurements to be carried out on lightweight, low power hardware that can easily be obtained in many parts of the region.

5.2 // Community participation

Several Internet Society Chapters are currently contributing to this project by hosting probes, which is helping to increase the number of measurement vantage points in Africa. Currently, the Madagascar, Benin, Tunisia, and Ethiopia Chapters are actively engaged in setting up measurement probes and have active infrastructure. We are already collecting data in Kenya and Mauritius on infrastructure that has been deployed by Internet Society staff (Kenya) and AFRINIC staff (Mauritius).

5.3 // Pilot phase

We have already started with increasing the number of Internet measurement vantage points (i.e., the MIRA Pods) in Africa by supplying measurement infrastructure and supporting deployment. We are already collecting—or preparing to collect—metrics on throughput, round trip time (RTT), and latency measurements in Benin, Burkina Faso, Congo DRC, Kenya, Madagascar, Nigeria, Tunisia, Rwanda and South Africa. More countries will be added as soon as suitable vantage points are identified.

5.4 // Partners

This project will be carried out in partnership with different academic institutions. Below is a list of principal investigators and institutions for the different work packages.

1. Amreesh Phokeer, AFRINIC and Kevin G. Chege, Internet Society will be responsible for overall coordination of the activities mentioned in the different Work Packages.
2. Assane Gueye, CMU-Rwanda will be responsible to work on WP2 to model infrastructure resilience framework.
3. Ahmed Elmokashfi, Simula Research Lab will work on WP3 to measure and map the physical and logical topology.
4. Josiah Chavula, University of Cape Town will be responsible for WP3 and WP4, measuring network resilience in terms the physical and logical choke-points, as well as quality of service at network and application levels.

Note: *AFRINIC and the Internet Society will be managing the MIRA project and will therefore jointly oversee all of the work packages defined above, making sure the different collaborators meet the requirements defined in terms of deliverables and deadlines.*

6. Ethical Considerations

The project team will take all necessary precautions to ensure that no personally identifiable data is presented in the public domain. All datasets will be anonymized before processing to prevent leakage of confidential information. Only the required data will be collected. Meta-information about the source (such as IP and geolocation) will not be stored and processed. We intend to aggregate data at two main levels: ASN-level and country-level.

Any dataset containing limited personal data will be used for the purpose of this project only and will not be transferred to any third party and will be discarded at the end of this project.

Furthermore, the MIRA Pods store data in JSON format which only contains the measurement data and not any personally identifying data. This is the data that will be processed and displayed on the MIRA dashboard within Internet Society Pulse for visualization.

7. Summary

Assessing the resilience of the Internet is an important activity in order to determine how Internet access and experiences can be improved. Contact information and information related to the project in general will be made available on the Internet Society[25] and AFRINIC[26] website and on Internet Society Pulse[27]. For details about the MIRA project and the measurement infrastructure, visit https://github.com/mira-project/mira/wiki.

[25] https://internetsociety.org
[26] https://afrinic.net/research/studies/mira
[27] https://pulse.internetsociety.org

References

[1] J. Rydzak, M. Karanja, and N. Opiyo, "Internet Shutdowns in Africa| Dissent Does Not Die in Darkness: Network Shutdowns and Collective Action in African Countries," *Int. J. Commun.*, vol. 14, p. 24, 2020.

[2] T. Freyburg and L. Garbe, "Blocking the bottleneck: Internet shutdowns and ownership at election times in sub-Saharan Africa," *Int. J. Commun.*, vol. 12, pp. 3896–3916, 2018.

[3] R. Kathuria, M. Kedia, G. Varma, K. Bagchi, and R. Sekhani, "The anatomy of an Internet blackout: measuring the economic impact of Internet shutdowns in India," 2018.

[4] R. Fanou, F. Valera, and A. Dhamdhere, "Investigating the Causes of Congestion on the African IXP substrate," in *Proceedings of the 2017 Internet Measurement Conference*, 2017, pp. 57–63.

[5] A. Formoso, J. Chavula, A. Phokeer, A. Sathiaseelan, and G. Tyson, "Deep diving into africa's inter-country latencies," in *IEEE INFOCOM 2018-IEEE Conference on Computer Communications*, 2018, pp. 2231–2239.

[6] A. Gupta, M. Calder, N. Feamster, M. Chetty, E. Calandro, and E. Katz-Bassett, "Peering at the Internet's frontier: A first look at isp interconnectivity in Africa," Lect. Notes *Comput. Sci.* (including Subser. Lect. Notes *Artif. Intell. Lect. Notes Bioinformatics*), vol. 8362 LNCS, pp. 204–213, 2014.

[7] M. Isah, A. Phokeer, J. Chavula, A. Elmokashfi, and A. S. Asrese, "State of Internet measurement in Africa-A survey," in *International Conference on e-Infrastructure and e-Services for Developing Countries*, 2019, pp. 121–139.

[8] R. Cohen, K. Erez, D. Ben-Avraham, and S. Havlin, "Resilience of the Internet to random breakdowns," *Phys. Rev. Lett.*, vol. 85, no. 21, p. 4626, 2000.

[9] M. Omer, R. Nilchiani, and A. Mostashari, "Measuring the resilience of the global Internet infrastructure system," in *2009 3rd Annual IEEE Systems Conference*, 2009, pp. 156–162.

[10] J. P. Rohrer, A. Jabbar, and J. P. G. Sterbenz, "Path diversification for future Internet end-to-end resilience and survivability," *Telecommun. Syst.*, vol. 56, no. 1, pp. 49–67, 2014.

[11] J. P. G. Sterbenz, E. K. Çetinkaya, M. A. Hameed, A. Jabbar, S. Qian, and J. P. Rohrer, "Evaluation of network resilience, survivability, and disruption tolerance: analysis, topology generation, simulation, and experimentation," *Telecommun. Syst.*, vol. 52, no. 2, pp. 705–736, 2013.

[12] J. P. G. Sterbenz et al., "Resilience and survivability in communication networks: Strategies, principles, and survey of disciplines," *Comput. Networks*, vol. 54, no. 8, pp. 1245–1265, 2010.

[13] J. Wu, Y. Zhang, Z. M. Mao, and K. G. Shin, "Internet routing resilience to failures: analysis and implications," in *Proceedings of the 2007 ACM CoNEXT conference*, 2007, pp. 1–12.

[14] J. P. G. Sterbenz et al., "Resilience and Survivability in Communication Networks: Strategies, Principles, and Survey of Disciplines," *Comput. Networks Spec. Issue Resilient Surviv. Networks*, vol. 54, no. 8, pp. 1245–1265, 2010.

[15] J. P. Rohrer and J. P. G. Sterbenz, "Predicting topology survivability using path diversity," in *2011 3rd International Congress on Ultra Modern Telecommunications and Control Systems and Workshops (ICUMT)*, 2011, pp. 1–7.

[16] R. Teixeira, K. Marzullo, S. Savage, and G. M. Voelker, "Characterizing and measuring path diversity of Internet topologies," *ACM SIGMETRICS Perform. Eval. Rev.*, vol. 31, no. 1, pp. 304–305, 2003.

[17] D. Baltrunas, A. Elmokashfi, and A. Kvalbein, "Measuring the reliability of mobile broadband networks," in *Proceedings of the 2014 conference on Internet measurement conference*, 2014, pp. 45–58.

[18] A. Csoma, A. Gulyás, and L. Toka, "On measuring the geographic diversity of Internet routes," *IEEE Commun. Mag.*, vol. 55, no. 5, pp. 192–197, 2017.

[19] K. G. Leyba, B. Edwards, C. Freeman, J. R. Crandall, and S. Forrest, "Borders and gateways: measuring and analyzing national as chokepoints," in *Proceedings of the 2nd ACM SIGCAS Conference on Computing and Sustainable Societies*, 2019, pp. 184–194.

[20] R. Fontugne, A. Shah, and E. Aben, "As hegemony: A robust metric for as centrality," in *Proceedings of the SIGCOMM Posters and Demos*, 2017, pp. 48–50.

Appendix 1

Category	Data item	Data type	Purpose/Research Questions
ISP Resilience -	Link resilience	Secondary	• How resilient an ISP is in terms of upstream connectivity?
	QoS/QoE	Primary	• What is the quality of the link (performance, uptime, reliability)? • Are end-users experimenting the same level of QoE?
	DNS Resilience	Primary	• Are the ISPs providing a resilient DNS resolver service? • Public DNS information
Critical infrastructure resilience	Cable	Secondary	• Are there chokepoints in the connectivity at the physical level? • Is there any concentration (business or geographical) on the landing ports? • Is there any concentration on the cable service provider?
	Power ecosystem	Secondary	• How resilient is the power
	ccTLD	Primary	• Number of name servers? • Location of name servers? • DNSSEC
Market Resilience	Chokepoint potential	Secondary Primary	• Do we see concentration towards a small group of upstream?
	Traffic Localization	Primary Secondary	• % of AS are peering at the IX? - % of AS exchanging traffic? • Number of IXes in a country • Amount of local popular content hosted in-country
	Affordability	Secondary	• How affordable is access to Internet connectivity

Appendix 2—2021 Plan

WP/Task	Actors	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec
WP 1. Measurement Infrastructure													
1.1 Measurement Pods	AFRINIC/ISOC	//	//	//									
1.2 Pods orchestration	AFRINIC/ISOC	//	//	//									
1.3 Data pipeline & storage	AFRINIC/ISOC		//	//									
WP 2. Internet Resilience Framework													
2.1 Theoretical framework	CMU/AFRINIC/ISOC		//	//									
2.2 Benchmarking/Ground truth	CMU/AFRINIC/ISOC			//	//								
2.3 Analytics pipeline	CMU/AFRINIC/ISOC			//	//	//							
WP 3. Resilience Analytics													
3.1 Physical topology analysis	UCT, Simula/AFRINIC/ISOC	//	//										
3.2 Logical topology analysis	UCT, Simula/AFRINIC/ISOC		//	//									
3.3 Mapping logical & physical topology	UCT, Simula/AFRINIC/ISOC			//	//								
3.4 Chokepoint potential/AS Hegemony	AFRINIC/ISOC				//	//							
WP 4. Network Performance													
4.1 Access performance	UCT, AFRINIC/ISOC					//	//	//					
4.2 Infrastructure peering	UCT, AFRINIC/ISOC						//	//	//				
4.3 Understanding web usage	UCT, AFRINIC/ISOC							//	//	//			
4.4 Network and web interference	UCT, AFRINIC/ISOC								//	//	//		
WP 5. Mobile broadband resilience													
5.1 Develop mobile measurements tools	TBD												
5.2 Usability of mobile broadband	TBD												
5.3 Critical services (QoE)	TBD												
WP 6. DNS Infrastructure resilience													
6.1 ccTLD robustness	AfTLD/ISOC												
6.2 DNS performance	AfTLD/ISOC												
WP 7. MIRA Dashboard													
7.1 API for external usage	AFRINIC/ISOC			//	//	//	//						
7.2 Visualization dashboard	AFRINIC/ISOC					//	//	//					