Internet Society Pulse

Internet Resilience Index Methodology

2021

October 2021 v1.0

Table of Contents

Introduction

About the Index

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 are on a level playing field with regards to a resilient Internet infrastructure. Many low-income countries usually have under-provisioned networks and cable infrastructure, or they lack redundant interconnection systems. In these countries (or regions), the likelihood of Internet outages occurring is much higher than in other countries.

Measuring Internet resilience is not an easy task as there are several building blocks underpinning the Internet's complex infrastructure. Additionally, the Internet landscape varies considerably around the world and to be able to objectively compare countries between one another - on a common ground - there needs to be an objective set of metrics that tracks and records the different components that contribute to the resiliency of the Internet.

To achieve this task, the Internet Society is introducing the Internet Resilience Index (IRI). This document outlines the approach used to build the Index, the selection of indicators and the underlying data sources, the weighting scheme, and the aggregation method used.

The Four Pillars of a Resilient Internet Ecosystem

In order to grasp the multi-faceted nature of the Internet, the Index is built on four main pillars, which together contribute to the smooth operation of the Internet. The pillars are:

1. **Infrastructure**: The existence and availability of physical infrastructure that provides Internet connectivity.
2. **Performance**: The ability of the network to provide end-users with seamless and reliable access to Internet services.
3. **Security**: The ability of the network to resist intentional or unintentional disruptions through the adoption of security technologies and best practices.
4. **Market Readiness**: the ability of the market to self-regulate and provide affordable prices to end-users by maintaining a diverse and competitive market.

The Internet Society IRI is built using existing best practices according to the EU-JRC and the OECD Handbook on Constructing Composite Indicators[1] and uses the same methodology as currently existing

[1] https://knowledge4policy.ec.europa.eu/publication/handbook-constructing-composite-indicators-methodology-user-guide-0_en

indices such as the GSMA Mobile Connectivity Index[2], the Facebook/EIU Inclusive Internet Index[3] and the Web Foundation Web Index[4].

Data Sourcing

Selecting Indicators

Building a robust composite indicator requires careful selection of the underlying indicators. To date, there are no direct and readily available metrics that provide information about the resilience of a network or a country. In the Internet Society IRI model, the indicators selected are reflective of a specific aspect of resilience that needs to be quantified. The OECD and EU-JRC handbook provides some guidance on the main characteristics to consider when selecting the indicators. In essence, they should be accurate, timely, and should cover as many countries as possible. Additionally, the Internet Society IRI relies exclusively on quantitative indicators as opposed to qualitative ones such as *perception of Service Quality*. This is to ensure that there is an objective set of metrics that can be used to make comparisons between countries.

Selection Criteria

The following criteria were used when selecting the datasets:

- **Relevance**: The indicator should work towards showing an increase or decline in the resilience of the Internet in a selected country.
- **Accuracy**: The indicator should correctly estimate or describe the quantities or characteristics they are designed to measure.
- **Coverage**: The data should cover as many countries as possible, as the Index is intended to be a global index. An indicator is not included if there is missing data on more than 25% of countries in the Index.
- **Freshness**: any dataset should be at most two years old. Some datasets such as performance or network coverage should be recent. Some other datasets such as number of exits points do not change considerably over years, so it is acceptable to use a dataset which is a year or two old.
- **Continuity**: In order to objectively compare the index over the years, it is important to work with a stable list of indicators, which will provide data consistently over time.

[2] https://www.mobileconnectivityindex.com/
[3] https://theinclusiveinternet.eiu.com/
[4] https://thewebindex.org/

Types of Indicators

There are three main types of indicators that have been used to calculate the Internet Society IRI:

1. **Direct indicator:** A direct indicator is a direct measure of an aspect of resilience e.g., percentage of HTTPS adoption, latency, bandwidth, etc. They have a specific unit of measurement, and the raw value can be on different scales depending on what is being measured.
2. **Composite indicator:** A composite indicator provides a score, which itself has been derived from multiple other variables. Examples are the MANRS score, EGDI index, Market Concentration, etc. The scale of a composite indicator is usually between 0 and 100.
3. **Proxy indicator:** A proxy is used where it is difficult to find a specific metric to measure an aspect of resilience. Proxies can be either direct or composite indicators. For example, IRI uses "Number of IXPs" and "Number of data centers", together to quantify the robustness of the local infrastructure.

Orientation of Indicators

An indicator can either be positive or a negative. In the Internet Society IRI model, both positive and negative indicators are used either individually or in combination with other indicators to depict a notion of resilience. An example of a positive indicator is "Number of secure Internet servers" as the higher the number the more secure will be the network. On the contrary, "% of spam infections" is a negative indicator, as the higher the percentage, the less secure the underlying networks are.

Details on Some Indicators

Exit Points

The number of exit points, or international gateways, is an important indicator of the resilience of a country in terms of infrastructure. For example, in the event of a submarine cable break, the Internet traffic can be redirected to a neighboring country with terrestrial fiber. Exit points are counted by looking at the number of cable landing stations[5] and/or terrestrial cross-border fiber connection points[6].

[5] https://submarinecablemap.com/

[6] http://www.africabandwidthmaps.com/

Network Performance

The data about bandwidth and latency is collected from the quarterly Ookla Speedtest open datasets[7]. It contains measurements about fixed and mobile network performance around the world. The median download, upload and latency values are calculated by country.

Peering Efficiency

The Peering Efficiency score of a country is calculated by taking the sum of all unique and local[8] (Autonomous Systems) ASes peering at an IXP in a country and dividing it by the number of allocated ASes in a country. PeeringDB[9] or Packet Clearing House (PCH)[10] provides the number of unique local peers and the RIRs delegated file[11] provides the total number of allocated ASes by country.

$$PE_c = \frac{\sum P_i}{A}$$

Where:

$$PE_c = Peering\ Efficiency\ of\ country\ c$$

$$P_i = Local\ ASes\ peering\ at\ IXP\ i$$

$$A = Number\ of\ allocated\ ASes\ for\ country\ c$$

Market Concentration

The Internet Society IRI uses the Herfindahl-Hirschman Index (HHI) to calculate the market concentration score. APNIC ASPOP statistics[12] provide market share information by AS and by country. The HHI has a range between 0 and 10,000, where 0 means no concentration (a competitive market) and 10,000 means only one ASN is present i.e., with 100% market share.

$$HHI_c = s_1^2 + s_2^2 + s_3^2 + \cdots s_n^2$$

Where:

$$HHI_c = HHI\ index\ of\ country\ c$$
$$s_n = market\ share\ (\%)\ of\ ASN_n\ of\ country\ c$$

[7] https://registry.opendata.aws/speedtest-global-performance/

[8] Both local and foreign ASes (e.g., CDNs, Tier-1 operators) peering at an IXP. For this calculation, only local ASes are considered.

[9] https://www.peeringdb.com/

[10] https://www.pch.net/ixp/dir

[11] https://ftp.ripe.net/pub/stats/ripencc/nro-stats/latest/nro-delegated-stats

[12] https://stats.labs.apnic.net/cgi-bin/aspop

AS Hegemony

Network centralization is an important element to measure as it indicates the extent to which the relationships of a given network are concentrated on a single network or group of networks. At a country-level, there are specific network operators providing international access and the more diverse the number of upstream Internet providers, the more resilient the country is in terms of network dependency.

The notion of network dependency can be proxied using AS Hegemony[13] which is a score given to a network to quantify its centrality as observed by BGP monitors. AS hegemony ranges between 0 and 1 and can be interpreted as the average fraction of paths crossing a node. The higher the AS Hegemony score, the higher the dependency on that specific network.

Each network in a country has an AS Hegemony score, based on how central it is for all other eyeball networks. In order to calculate the inequality in the network dependency distribution at a country-level, we use the GINI coefficient[14] of inequality. In a perfectly equal scenario (GINI = 0), all networks would have the same AS Hegemony score. GINI=1 means perfect inequality.

Examples:

- Country C has three networks AS X (10% coverage), AS Y (50% coverage), AS Z (40% coverage). The GINI coefficient of country C is G(5, 10, 85) = 0.533 (High GINI)
- Country D has three networks AS X (33% coverage), AS Y (33% coverage), AS Z (34% coverage). The GINI coefficient of country D is G(33, 33 , 34) = 0.07 (Low GINI)

[13] https://www.iij-ii.co.jp/en/members/romain/pdf/romain_sigcomm2017.pdf

[14] https://en.wikipedia.org/wiki/Gini_coefficient

List of indicators

Table 1 shows the list of indicators, the unit of measure and the source of the information.

Indicator	Description	Unit	Source
Exit points (Gateways)	Exit points (terrestrial or submarine) used for upstream connectivity.	Number of physical exit points	Africa Bandwidth Maps
10-km Fiber reach	Percentage of the population within 10 km of a fiber connection point	% of population	ITU
Network Coverage	Mobile Network coverage includes 2G/3G/4G with a composite score provided by the GSMA	Score (0 - 100)	GSMA
Spectrum allocation	Spectrum allocation (composite score)	Score (0 - 100)	GSMA
Number of IXPs	Number of IXPs	# of IXPs per 10 million population	PCH PeeringDB
Getting electricity	Getting Electricity Index incorporates the quality of power supply	Score (0 – 100)	World Bank
Data centers	Number of datacenters	# of datacenters per 10 million population	Datacentermap
Mobile/Fixed Latency	Median latency observed to the nearest Ookla server	ms	Ookla
Mobile/Fixed Upload speed	Median upload throughput measured to the nearest Ookla server	Mbps	Ookla
Mobile/Fixed Download speed	Median download throughput measured to the nearest Ookla server	Mbps	Ookla
IPv6	IPv6 Enabled end users	% of IPv6 adoption	Akamai, Facebook, Google, APNIC
HTTPS	Pageload using HTTPS	% of pageload on HTTPS	Mozilla
DNSSEC Validation	Users validating DNSSEC	% of users validating DNSSEC	APNIC
DNSSEC Adoption	Is the ccTLD signed?	0 or 1	ICANN

MANRS Readiness	MANRS score (Filtering, Global Coordination, IRR, RPKI)	Score (0 – 100)	MANRS Observatory
Secure Internet Servers	Number of Secure Internet Servers detected on the country's network	# of Secure servers per 1000 population	World Bank
Global Cybersecurity Index	Global Cybersecurity Index (Composite score)	Score (0 – 100)	ITU
DDOS Potential	Potential DDOS threat a country represents	Percentage	Cybergreen
Spam infections	Percentage of allocation that is listed in CBL (Composite Blocking List)	Percentage	Spamhaus
Affordability	Average of affordability for fixed and mobile broadband	% of GNI per capita	ITU/A4AI
Market concentration	Herfindahl-Hirschman Index (HHI) calculates the market concentration based on market share information per network.	Score (0 – 10000)	APNIC
Network Centralization	GINI Coefficient is used to calculate the inequality in the dependency on specific network for upstream connectivity.	Score (0 – 1)	IIJ
Peering efficiency	Ratio of ASes peering at IXPs vs allocated ASes in a country	Percentage	PCH PeeringDB
Domain count	Domains registered by ccTLD	# of domains per ccTLD per 1000 pop.	DomainTools
EGDI	E-Government Development Index	Index (0 – 100)	UN

Table 1. List of indicators

Data Processing

Raw data comes in different forms and shapes and usually comes with several artifacts - some datasets are normally distributed, while some others are skewed. Before running any calculation or aggregation it needs to be imputed for missing data and treated for outliers.

Missing Data

The following techniques have been used to impute missing data:

Indicator	Technique	Details
Getting Electricity	Interpolation	Used data from previous years.
Affordability	Mean imputation	The affordability value is calculated by taking the mean between the fixed and mobile values. If one value is missing, no average is calculated.
Latency, download and upload speed	Substitution	M-Lab speedtest dataset was used as a replacement.
Data centers	Substitution	Infrapedia.com was used as an alternate source of data.
Other indicators	Substitution	Replacement by data from a similar country based on GPD per capita.

Table 2. Data imputation

Re-scaling and Treating Outliers

The scales used by the indicators are also different e.g., latency can range between 0 – 500ms, while domain count for a ccTLD can range between 0 – 2,000,000. It is important to scale the data so that indicators are comparable to one another, and to avoid the issue of the size of the country (i.e., larger countries in terms of population or GDP tend to have more networks, IXPs, datacenters, etc.).

On the other hand, outliers have the tendency to skew the data and can therefore have an impact on the overall score calculation, especially that Internet Society IRI uses the min-max normalization (see section on Min-Max Normalization below) method to scale the data. If an indicator has a very high or very low value, this will be reflected in the min-max calculation.

The following treatment has been applied to the specific indicators of the framework:

1. Normalization by population size: Number of data centers, number of IXPs, Secure Internet Servers, Number of domains

2. Log transformation[15]: HHI Index, Secure Internet Servers, fixed/mobile upload speed, fixed/mobile download speed and fixed/mobile latency.

After scaling and transforming the above indicators, we run a check on the skewness and kurtosis values on the remaining indicators. For those having a skewness > 2 and kurtosis > 3.5 (general threshold for outlier detection), IRI makes use of the IQR (Interquartile Range: Q3 -Q1) method to trim down outliers. The following rules are applied:

- Any value greater than Q3 + 1.5*IQR, is replace by Q3 (75th percentile)
- Any value less than Q1 – 1.5*IQR, is replaced by Q1 (25th percentile)

Min-Max Normalization

The next step, after cleaning and transforming the data, is normalization. Normalization is important because indicators are collected using different unit of measurements (percentage, ms, Mbps, count, etc.). It is therefore important to rebase them to a common unit such as into 0 to 100 scale, where 100 usually refers to the strongest and 0 to the weakest value.

The method chosen was the *min-max* normalization, it is a common technique used by multiple known indices[16][17] and as opposed to other techniques such as ranking and categorical scales, *min-max* keeps the interval between the countries consistent.

Below are the formula Internet Society IRI uses to calculate the value of an indicator depending on whether it is positive or negative:

Positive indicator	Negative indicator
$I_{k,c} = \frac{x_{k,c} - Min(x_k)}{Max(x_k) - Min(x_k)}$	$I_{k,c} = 1 - \frac{x_{k,c} - Min(x_k)}{Min(x_k) - Min(x_k)}$

where "x" refers to the raw value of the indicator "k" of country "c" and "I" refers to the normalized value. $Max/Min(x_k)$ refers to the min/max of indicator "k" for all countries.

[15] A Logarithmic transformation is useful to treat skewed datasets and to discard extreme values. Not only it scales the data, but it has the advantage of handling outliers in the dataset. Log transformation preserves the differences between the values.

[16] EIU Internet Inclusive Index - https://theinclusiveinternet.eiu.com/

[17] GSMA Mobile Connectivity Index - https://www.mobileconnectivityindex.com/

Positive indicators contribute towards increasing an index, negative indicators contribute to decrease the score, which is why we take the delta $(1 - I_{k,c})$.

Finally, we chose not to use the *z-score* standardization[18] technique as not all the indicators followed a normal distribution.

Weighting and Aggregation

Assigning Weights

There are two main ways to aggregate the normalized indicators into a final score: (1) using an ad-hoc weighting scheme (2) using statistical (optimization) techniques. Internet Society IRI uses a weighting scheme as it is the simplest technique between the two and relies on input that the Internet Society gathered through survey and discussions with subject matter experts.

During the weighting process, the importance of the indicator was also considered using a "lifecycle" approach. For example, for the Infrastructure pillar, the following weights were assigned to the underlying dimensions: Cable Ecosystem (40%), Mobile Connectivity (30%) and Enabling infrastructure (30%). Higher importance was given to *Cable Ecosystem* as it is a prerequisite for a functional Internet.

In the Internet Society IRI model, the indicators are grouped into different dimensions, and the dimensions into pillars, which on their own provide quantitative measures of a specific aspect of Internet resilience. Below is a table showing the indicators, dimensions and pillars and their associated weights, used for the calculation of the Internet Society IRI in 2021.

The weights will be revisited on an annual basis.

Pillar	Weight (%)	Dimension	Weight (%)	Indicator	Weight (%)
Infrastructure	25	Cable ecosystem	40	Exit points (Gateways)	50
				10-km Fiber reach	50
			30	Network Coverage	70

[18] This technique standardizes around the mean value and ranges between 0 and 1.

		Mobile connectivity		Spectrum allocation	30
		Enabling infrastructure	30	Number of IXPs	20
				Power availability	40
				Datacenters	30
Performance	25	Fixed networks	40	Latency	30
				Upload	30
				Download	40
		Mobile networks	60	Latency	30
				Upload	30
				Download	40
Enabling technologies and security	25	Enabling technologies	20	IPv6	30
				HTTPS	70
		DNS ecosystem	30	DNSSEC Validation	50
				DNSSEC Adoption	50
		Routing hygiene	30	MANRS Readiness	100
		Security threat	20	Secure Internet Servers	30
				Global Cybersecurity Index	30
				DDOS Potential	20
				Spam infections	20
	25	Market structure	50	Affordability	40

Local ecosystem & Market readiness				Market concentration	30
				AS Hegemony	30
		Traffic Localization	50	Peering efficiency	40
				Domain count	20
				Local content	20
				EGDI	20

Table 3. Indicators, dimensions and pillars and associated weights for year 2021

Aggregation

IRI uses a weighted sum formula at each level (indicator, dimension and pillar) to aggregate the data into a composite score. The following formula was used:

$$IRI_c = \sum_{i}^{n} (w_i * P_{i,c})$$

Where:

$$P_{i,c} = \sum_{i}^{n} (w_i * D_{i,c})$$

And where:

$$D_{i,c} = \sum_{i}^{n} (w_i * I_{i,c})$$

In simple terms, the final index IRI_c of country "c" is the sum of the weighted pillars "P_i". A pillar is the weighted sum of the underlying dimensions "D_i" and a dimension is the weighted sum of the indicators "I_i" all of country "c".

Interactive Dashboard

Since the weights were assigned following a qualitative methodology, it is understood that different users may have different opinions on the weighting scheme. For example, an end-user might be more interested in *Performance* (thereby assigning a higher weight) than *Security*. Similarly, a regulator might be more inclined to add more weight to the *Market Structure* dimension than to the *Traffic Localization* dimension.

Therefore, our approach is to allow users to adjust the weights to reflect on-the-ground realities. By zeroing the weight, a user can even remove one indicator, dimension or pillar from the calculation of the final index. This will be achieved through an interactive dashboard, which is currently under development and will be available in Q1 2022.

Feedback

For any questions, comments and feedback on the IRI, please contact the Internet Society Pulse team (pulse@isoc.org).

Acknowledgements

The Internet Society would like to thank the following contributors for their valuable input in the conception of the Internet Society Internet Resilience Index (IRI). Amreesh Phokeer (Internet Society), Kevin Chege (Internet Society), Assane Gueye (Carnegie Mellon University-Africa), Josiah Chavula (University of Cape Town), and Ahmed Elmokashfi (Simula Research Lab).