

Internet eXchange Federation Project Funding Proposal

The Internet eXchange Point Database (IXPDB)

Executive Summary

Internet Exchange Points (IXPs) are network facilities that enable the interconnection of three or more independent networks, primarily for the purpose of facilitating the exchange of Internet traffic. IXPs are considered core Internet infrastructure and carry a growing percentage of the world's Internet traffic.

While there are a number of disparate websites and resources, which contain various data on IXPs, none are globally authoritative and fully comprehensive. This makes it difficult for network operators, researchers and other interested parties to make informed decisions related to the global Internet ecosystem.

We, as the Internet eXchange Federation (IX-F), aim to develop a system that will fill this gap. It will acquire data directly from IXPs through automation; aggregate and promote useful third-party data sources; and provide valuable insight, reporting, data export and visualisation functionality. We aim to work closely with the Internet technical community throughout the project and will host workshops that help to promote use of the system and encourage the application of best common practices (BCPs).

About the Internet eXchange Federation (IX-F)

The Internet eXchange Federation (www.ix-f.net) is a transnational organization for coordinating the efforts of Internet eXchange Point Associations (IXPAs). Our membership includes the African Internet Exchange Point Association (www.af-ix.net), the Asia Pacific Internet Exchange Point Association (www.apix.asia), the European Internet Exchange Point Association (www.euro-ix.net), and the Latin America, and Caribbean Internet Exchange Point Association (www.lac-ix.net). Collectively, these IXPAs represent over 145 IXPs across 6 continents.

Introduction to the IXP Database (IXPDB)

The Internet eXchange Point Database (IXPDB) aims to be the authoritative, comprehensive, public source of data related to IXPs. The IXPDB collects data directly from IXPs through a recurring automated process that integrates with common IXP measurement and management software. It also integrates data from second and third party sources in order to provide a comprehensive and corroborated view of the global interconnection landscape. The combined data can be viewed, analyzed, and exported via a powerful web-based interface and software API.

Database Specifications

First-party data collected directly and automatically from IXPs will include:

- Essential information on IXP providers, platforms, and point of presence (PoP) locations.
- Aggregate traffic data with historical values.
- Detailed information on the networks present at each IXP including their ASN, join date, switch, switch port, port pricing, port capacities, VLANs, peering policy, route server usage, max prefix setting, prefix announcements, and contact data.
- Hardware vendor data for IXP switches and network participant routers.
- Route server hardware, software, and version information.
- IXP IP address allocations and assignments for v4 and v6.

Second and third party data sources will include:

- PeeringDB: Essential information on IXP facilities (e.g. datacenters) and additional information on connected networks including network type (e.g. content, NSP), public peering locations, private peering locations, peering policy, ratios, and contact data.
- PCH Routing Data: Detailed technical information about connected networks and their prefix announcements as seen by PCH route collectors in over 140 IXPs.
- RouteViews, WHOIS, RIPE Atlas, and CAIDA.

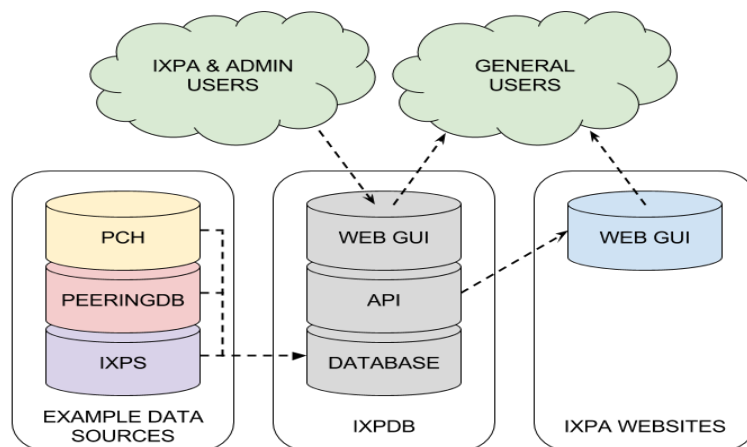
Example use cases and queries:

- An IXPA wants to conduct analysis on its' membership in order to generate reports, identify gaps, and promote investment.
 - Which IXPs in my region are growing at the fastest rate? What are the trends?
 - What is the total amount of aggregate traffic in my region? What are the trends?
 - How many networks are connected to IXPs in my region in aggregate, directly and indirectly, and how has that number changed over time?
 - What types of networks are connecting to IXPs in my region, directly and indirectly, and how has that diversity changed over time?
 - What is the most popular hardware vendor in my region, and how has this changed over time?
- A network operator or CDN wants to leverage peering as they expand into new markets
 - Which IXPs have a high density of ASNs and prefixes that I need to reach?
- An enterprise wants to evaluate which network operators to use for transit.
 - Which network operators available to me are peering with other networks that I need efficient connectivity to in various markets?
- A researcher wants to identify gaps in a region in order to better allocate development resources or measure the efficacy of past development projects.
 - Can I identify any correlations between previous development projects and improvements in the interconnection ecosystem?
 - Which sub-regions in Africa are performing best across various metrics?
 - Which regions or markets are underperforming that might be improved by additional development projects or resource allocations?
- A government entity wants to check the status of the Internet ecosystem in order to substantiate the efficacy of national or regional policy implementations.

- Have more international networks connected to IXPs in a country or region since the implementation of new policies and regulations?
- Have large-scale government datacenter deployments been successful, and have they had a positive impact on the interconnection and content ecosystems in their countries?
- How does my country or region's interconnection ecosystem compare to others across various metrics.
- An IXP operator wants to benchmark itself against other IXPs in the region and identify trends that can inform its growth strategy.
 - What new networks or network prefixes are starting to appear in my region?
 - Which networks in my country or region are not connected to my exchange?
 - Where do ASNs and prefixes that originate from my region appear elsewhere in the world? What are the trends? Should I offer a remote peering service to IXPs?

System Architecture

The following abstract diagram illustrates the structure of the system and the logical flow of information between the various components and users:



Summary of information flow:

1. IXPA and IXPDB Administrators create entries for IXPs in the IXPDB system
2. The IXPDB system begins automatically collecting information about those IXPs from various sources (e.g. IXPs, PeeringDB, and PCH)
3. IXPDB presents collected data to general users via an intuitive web-based interface located on the IXPDB website.
4. IXPDB data is also be exported to IXPA websites using a well documented API

Development Plan

The software development is broken down into baseline data and features and a longer-term plan for introducing and refining sophisticated market analysis features.

Objectives for the first 12 months:

- Develop and release version 1 of the IXPDB system with automated data collection from all public IXP Manager instances and PeeringDB; granular data export functionality via an API; and essential web-based administration features.
- Automated first-party data collection from at least 40 IXPs; 10 from each IXPA.
- Conduct one training and integration workshop within each IXPA region which solicits feedback in order to inform future system development work
- Develop a project plan for the fully-fledged web-based GUI.

Objectives for the following 24 months:

- Automated first-party data collection from at least 80 IXPs; 20 from each IXPA.
- Integration of additional second and third-party data sources
- Additional training and integration workshops with an additional focus on utilization
- Develop online tutorials, videos, and documentation
- Improve the web-based user interface for better usability (e.g. multilingual support, additional data visualisations, enhanced query and report customization)

Project Team

Bijal Sanghani, Secretary General at Euro-IX, is the Project Lead with support from two dedicated representatives from each Internet Exchange Point Association (IXPA).

The software development work has been contracted to CZ.NIC, which operates the .CZ DNS registry. CZ.NIC has a long track record of developing open source software and infrastructure projects, including the BIRD routing daemon, Knot DNS, Knot Resolver, and TLD registry system FRED. Michael Horejsek of CZ.NIC is the Technical Project Manager responsible for the website and database development, he reports to Bijal Sanghani.

Jesse Sowell, currently a Cybersecurity Postdoctoral Fellow at Stanford's Center for International Security and Cooperation, is providing data architecture design and analytics prototyping for the IXPDB data architecture, IPP tools, infrastructure economics analyses, and attendant prototyping of data visualisation. In the coming year Jesse will be taking a faculty position where a key part of his team's research portfolio will be Internet infrastructure analyses, much of which will be built on IXPDB data.

The project team will also coordinate with the broader Internet technical community in order to obtain input and feedback throughout the project work.

Project Budget

The table below represents the estimated budget for the first 12 months:

Budget Item	Cost
Project Coordination	€40,000
Software Development	€72,000
Cloud Hosting Services	€3,000
Integration support / training / workshops for IXPs and IXPAs	€40,000
Total	€155,000

Table Breakdown

- **Project Coordination:** We would like to hire a dedicated project coordinator to support the development process and integration with IXPAs and IXPs. This role involves organizing workshops on system implementation, system utilization, best common practices, and documentation.
- **Software Development:** Our Technical Project Manager and his team at CZ.NIC will be doing the back-end and front-end development.
- **Integration support, training, and workshops:** Costs related to the development of training materials, supplies, and logistical costs (e.g. venue, travel, accommodation) for the workshops and training events.

We estimate that our second and third year budgets will have similar totals, with a shift in focus toward training (using a train-the-trainer approach) and implementation. We expect a significant reduction in software development costs in subsequent years.

Sponsorship Packages

In order to achieve our objectives, we are seeking support via:

1. Sponsorship: Make a financial contribution towards the cost of the project
2. Participation: Providing a dedicated person from your organisation to provide input during the project's software development and implementation

	Platinum	Gold	Silver	Bronze
	€25,000	€15,000	€10,000	€5,000
Two members of your organisation can attend all IXPA meetings in the year of sponsorship	X			
Be the first to test the new code and get on our beta test group	X	X		
One member of your organisation can join the development mailing list	X	X	X	
Sponsorship level size logo displayed on Sponsors page	X	X	X	X
Sponsorship level size logo displayed on marketing material	X	X	X	X
Receive end of year report on status of project and funds allocated	X	X	X	X

Want to sponsor us or want more information? Get in touch, secretariat@ix-f.net.