Pipes: a Peer-to-Peer Application Development Platform

Pipes is the working name for a new platform for peer-to-peer (P2P) application development.

Google told the world the future of computing is in the browser, and the world followed -- to Google's great benefit. In a browser-based, cloud-dominated world, everything is controlled by a few strong players; all of your data, even your ability to work, relies on large, remote data and compute *centers*. They are a single point of failure for the world: when AWS or Squarespace goes down, millions are lost by businesses worldwide. The underlying protocols of the internet were originally designed for peer-to-peer use, and it is an accident of history that such centralized models remain predominant. We now have the compute power to do the majority of the work we need with our own resources, and in such a world P2P is the future of computing. The Pipes platform will contribute by providing a general toolkit that enables developers to harness that distributed power, enabling users to take back control of their software and data.

Peer-to-peer applications have been successful in niche markets and domains like Bitcoin, Napster, and the Git revision control system. These applications demonstrate the potential power of P2P designs, but the inherent difficulty in P2P engineering accounts for the lack of more widespread adoption. Pipes provides an API that simplifies P2P development on top of a platform that implements the most difficult aspects of P2P communication and synchronization.

With the power and simplicity of the Pipes platform, more developers (and users) will be able to take advantage of the benefits of P2P applications: decentralization, data ownership and control, and anti-scaling (see below). The flexibility of the Pipes API will allow for new kinds of collaborative applications to developed, removing global reliance on major service providers and fostering innovation in how remote teams get work done.

DECENTRALIZATION AND DATA OWNERSHIP

Where are your family photos? Your contracts and correspondence, your movies and music, your spreadsheets and source code? We trust our most treasured files to centralized services -- should we? We need solutions that allow complete control over data privacy, storage, archival, and encryption. Institutions have gone so far as to develop their own internal collaboration tools *de novo*, rather than trust their data to a third party. P2P is the solution. P2P models enable finer-grained control over security and privacy, allowing developers and administrators to tailor security practices to specific situations. The Pipes platform will provide a provably secure, auditable system for deploying or developing such applications, with all data kept "on-prem" or restricted to VPNs. For small groups that work loosely together, encrypted data can be synchronized through intermediary "routing" nodes, so that remote collaborators can securely synchronize work without relying on *any* third-party data storage.

Applications developed and deployed on the Pipes platform will have an unprecedented ability to control the security requirements of application and data transfer. For some institutions -- including those in the EU under the GDPR -- how and where data is stored and archived is of critical importance. On the Pipes platform, you can control how, where, and under what requirements data is encrypted and shared. Strong cryptography is built in from the ground up, and our approach originates in our intimate knowledge of blockchain protocols. Data synchronization is both encrypted and verified, with an optional cryptographically verifiable audit trail for industries such as law or medicine in which that is important. Whether you're a close-knit group of a few innovative developers that want to share a git repository or a large organization synchronizing user data to mobile devices, the Pipes platform allows you clear and systematic control of how and where data is transferred between peers.

ANTI-SCALING

How much of your call time do you spend troubleshooting the call itself? We've all guessed at missed words as we try to figure out whether the problems are due to our own wifi or on the service provider's end. On our internal P2P collaboration tool, we use the crystal-clear Opus 48k codec for voice communications: we have no connection issues, and the audio quality is clearly superior to that of "standard" tools such as Slack, Skype, Zoom, and Facetime Audio. Why don't these bigger players provide better quality? It comes down to the difficulty -- and cost -- of scaling. And because of it we all suffer from the lack of quality, the slow connections, and the strange failures. When using a P2P platform, scaling is not a concern, because you only need to handle the size of your workgroup -- a matter of tens or hundreds of users, not millions. A $35 Raspberry Pi with a public IP has all you need to route and multiplex high-fidelity audio and video because it never has to handle too many peers at once. This not only improves streaming media quality, but opens up other collaborative design techniques that are simply unavailable in a massive-scale client-server model. Stateless servers and protocols, federated database clusters, load balancers, the C10k problem -- all of these are relics of a client-server model that has reached the limits of its utility, and all are non-issues when developing on a P2P network.

POWERFUL API

Much remote work relies on the basic infrastructure of collaboration: chat, voice, video, screen sharing, and file sharing. These tasks are not inherently hard, but nearly all of us rely on a few centralized service providers, and as a result suffer poor quality and user experience -- and an IT nightmare of integration. With the Pipes platform, you create tools that focus on the primary work -- in our case, developer tools -- and then integrate the basic services. These basic services are provided as a matter of course, as a simple API call; this way, collaboration tools can be focused on specific industries and markets, and better support the needs of users.

Under the hood, the Pipes platform manages:

- Efficient routing of messages and streaming media
- Redundant routing, ensuring fast packet delivery even if intermediary nodes lose connection
- Secure communications and private, pairwise, and group-level encryption
- Identity management and cryptographic identity verification
- Synchronization of data between nodes, including background sync of large (filesystem) data
- Offline work and re-synchronization when connected

This last point is important from a user perspective: you don't need to be connected to the network to get work done, only to synchronize data and exchange messages.

Pipes provides developers with a straightforward API for accessing this power. Our team has a long history of creating developer tools and APIs, and with Pipes we have achieved a balance between simplicity of development and expressive power. The downfall in the design of many distributed systems is attempting to support the most general case, which is extremely difficult; in our implementation we strike a balance that supports most use cases while maintaining simplicity at the API interfaces and providing lightning-fast routing and synchronization.

Secure and effective remote collaboration tools are increasingly vital, as spotlighted by the COVID pandemic. As the way we all work experiences this upheaval, many organizations are sacrificing data security and privacy in order to get work done. This is a motivating factor in our development of Pipes, and we've invested the the past six months in earnest R&D.

The Pipes platform and API are currently in Beta, and we're seeking partnerships with developers and organizations to help test and develop the API. Contact us (info@496code.com) if you're interested in developing with the API, investing in the Pipes project, or in having P2P collaboration tools developed and tailored to the needs of your organization or industry.