Infrastructure Services RFI

A Request for Information issued on 2023-02-07

IETF Executive Director
exec-director@ietf.org

About the IETF

The Internet Engineering Task Force (IETF) is the premiere Internet standards body creating open protocols to ensure that the global Internet is built on the highest-quality technical standards. These standards, shaped by rough consensus and informed by running code, are developed by a large volunteer community of leading engineering and technical experts from around the world. IETF processes are open and transparent, and IETF standards are freely available to anyone.

Standards and protocols developed at the IETF provide a core framework for today's online world. Everything from video conferencing, to email, to cloud storage is built on standards developed in the IETF community. In short, our work makes the Internet work.

www.ietf.org

Overview

The current contract for IETF IT infrastructure services is a black box contract - we specify the systems to be maintained along with a very basic SLA, and the provider is responsible for the underlying infrastructure on which those systems operate, including the system administration strategy. The provider has always been much more accommodating than that in responding to requests for changes inside the black box, but these requests were for a long time piecemeal and not part of an overall strategy.

There have been regular complaints from the community that our infrastructure is more prone to failure, more difficult to upgrade and more difficult to accommodate new services, compared to a modern best practice infrastructure. The IETF Administration LLC has consulted with the community on the best way to address this and worked with the community to develop a new operational strategy for how the infrastructure should be operated.

The IETF intends to issue one or more RFPs for a service provider to migrate and manage our infrastructure services in line with this new strategy, but before doing there are a number of decisions to be made on how those RFPs are to be structured. This document is a Request For Information to help the IETF make these decisions and issue high quality RFPs.

RFI Process

Please note the following:

- This is a Request for Information (RFI), it is not an RFP.

- We expect to issue RFPs for our infrastructure services in the first half of 2023 with migration to the new infrastructure complete by the end of 2023. Those RFPs will be open to all. Full details of our RFP process are available on our website.[1]

- You do not need to reply to this RFP to be eligible to respond to the follow-up RFPs.

- If you can answer our questions then that would be very helpful but you are under no obligation to. Responding to this RFI will not result in any favoured status for subsequent bidders.

- All responses to this RFI will be anonymous but not confidential. We may issue a report to our community summarising and quoting the feedback we

---

[1] https://www.ietf.org/about/administration/rfps-and-contracts/

received and we reserve the right to anonymously use specific wording from the responses in our RFPs if we feel that would enhance the process and can be done fairly to all bidders.

- Please do not use this as an opportunity to try to put in an early bid for our RFPs. We are likely to disqualify anyone who we think is attempting to game this process.

- Please send your responses to ietf-rfps@ietf.org.

Timeline

07 February 2023   This RFI Issued

26 February 2023   This RFI closes

New Strategy

This section documents the new strategy developed in consultation with the IETF community. While we have no plans to revise this strategy though respondents to this RFI are free to comment and recommend changes.

Goals

The IETF IT infrastructure should be operated to meet the following goals, each of which is described in more depth in the sections below. It is recognised that some of these goals overlap or are interdependent.

1. **Fit-for-purpose service availability**
2. **Fit-for-purpose service performance**
3. **Separated, cloud-first services**
4. **Automated, transparent and accessible infrastructure management**
5. **Secure and enduring services and data**
6. **Comprehensive service monitoring**

Fit-for-purpose service availability

The IETF is a global community with an uneven geographic spread, and the community uses the IETF infrastructure 24x7x365. During IETF meetings, both the triannual plenaries and the many interim meetings, demand on services is high and uninterrupted availability is expected. At other times, the activities are not so time critical that they cannot tolerate a delay of a few hours if planned and sufficient warning is given.

The IETF requires its IT infrastructure to support a fit-for-purpose service availability, which means:

- Minimal unplanned downtime.
- Infrastructure designed to eliminate planned downtime, recognising that some of the applications may still require that.
- Where planned downtime is required, then:
  - It must not be during or during the preparation phase of IETF meetings or other key events..
  - Should be able to be scheduled for any day of the week at any time, to meet IETF operational requirements.

Fit-for-purpose service performance

The demands on IETF services varies considerably with particular peaks in and around IETF meetings. Also, there are a number of services intended for users to pull large amounts of data (1-10 GB) in order to maintain local copies of large datasets. While service performance can be adversely affected by a bottleneck at a level of the stack above the IT infrastructure, the IT infrastructure should not become the bottleneck.

The IETF requires its IT infrastructure to provide a fit-for-purpose service performance, which means:

- Excellent infrastructure performance.
- All relevant and potentially relevant performance/utilisation data collected.
- An infrastructure that scales to match load, particularly during key events, with minimal manual intervention.
- An infrastructure that recognises the global nature of the IETF and ensures excellent performance to all end users.
- An evidence based approach used in setting all resource limits.
- A strategy for rapidly addressing performance bottlenecks.

Separated, cloud-first services

The IETF has a wide variety of services, a high rate of change/growth in services, increasing integrations between services, significant fluctuations in usage with an underlying increase, and a wide variety of deployment models.

The IETF requires its IT infrastructure to operate cloud-first and support IaaS, PaaS and SaaS, which means:

- All services run in the cloud on public cloud platforms.

- All cloud platforms to support orchestration, including automated deployment, scaling and management, except where the application itself cannot support this.
- All OS-level services to be containerised so that they are entirely separated at the filesystem/package level and any one can have any supporting package upgraded without any other service being affected.
- Any service can be customised, upgraded and migrated to another platform with either no dependency or only entirely unavoidable dependency, on any other service or component of any other service.
- All services to support operation behind Web Application Firewalls, Content Delivery Networks and other front end services.

Automated, transparent and accessible infrastructure management

The IETF has an active community of volunteers contributing to its systems development efforts and its meetings' NOC team, and it is likely that there will be a similar depth of community contribution to our IT infrastructure management, if this is enabled. In addition, the IETF has a small and growing team of in-house developers and key development contractors, who like much of the development community, are increasingly adopting devops practices.

The IETF requires the management of its infrastructure to be automated, transparent and accessible, which means:

- All build and configuration managed through an automated configuration management and deployment platform (e.g. Ansible).
- Automation scripts in a public Git repository.
- Credentials and other secret information used in automation scripts to be properly protected.
- A full test environment with the expectation that wherever possible, deployment involves first deploying to test, validating efficacy of changes, and then deploy to production.
- Where reasonable, it should be possible for any member of the IETF community to build a replica of any IETF service, with placeholder information available to replace any confidential information.
- Active response to community contributions both on mailing lists and through a formal change management system (e.g. GitHub Pull Requests).

Secure and enduring services and data

The IETF is vulnerable to the same threats as any other organisation and needs to mitigate those at many levels. The threat model for the IETF is unusual with more of an emphasis on data integrity and preserving the accuracy and availability of

historical data, than on protecting confidential information. Where the IETF does collect highly confidential information, such as for the NomCom process, every effort is made to compartmentalise that. Additionally, the IETF receives significant nuisance traffic (as a proportion of overall traffic).

The IETF requires its services and data to be secure and enduring, which means:

- An embedded risk-aware culture, with regular peer review and external audit of all strategies, processes and systems.
- Security-first network/service design and network/service management.
- A formal access control model with centralised observability.
- Clear compartmentalisation of confidential information.
- A patch management process that minimises the threat from unpatched systems.
- A backup and restore strategy that provides strong assurance of data integrity and high confidence of system rebuild.
- Active management of nuisance traffic.

Comprehensive service monitoring

As the IETF services evolve, they are becoming more integrated and more interdependent making the task of identifying the source of problems a more complex task and ultimately reliant on good data. The tools used for capturing, processing, analysing and presenting data have developed rapidly in recent years and within any team that relies on data there will be a wide range of tools used. Also, as noted above, the number of data consumers is likely to grow significantly as the IETF community begins to contribute to our IT infrastructure.

The IETF requires comprehensive, standards-based service monitoring, which means:

- Every part of the infrastructure is instrumented.
- Centralised collection of monitoring data to enable cross-service analysis.
- Controlled publication of monitoring data that maintains operational security while providing maximum access to the IETF community.
- Where possible, standards based data collection and distribution, and where proprietary APIs need to be used then these must be open and documented APIs.

Overview of Current Infrastructure

This is a basic overview of our infrastructure and we will not be providing any more information than this at this stage. When we issue RFPs they will include all the details necessary for someone to bid to operate our infrastructure.

Hosting

Most of our key services are hosted on one of three servers located in a Fremont, USA, datacenter managed by our current service provider. Two run Xenserver, the other XCP-ng but each only has a single OS installation (OpenSuse) into which all applications are deployed. Some applications run inside Docker containers.

The new strategy envisages all of this moving to the cloud with each application in its own container and OpenSuse replaced.

CDNs

We use CloudFlare as our CDN and are regularly expanding our usage. CloudFlare provide their services to us free of charge under their Galileo project[2].

The new strategy envisages us using CDNs as appropriate.

Applications

Our applications are a mix of custom software and off-the-shelf software that has been customised for our environment. The key self-hosted applications in use are:

- Datatracker. An internally developed workflow management system written in Python/Django using MySQL and PostgreSQL for databases with MySQL being phased out. This also provides an OAuth server for SSO across our applications.
- www.ietf.org. A Wagtail website
- Interactive tools at author-tools.ietf.org. A web interface to a set of custom tools, mostly in Python..
- Satellite websites (rfc-editor.org, iab.org, iesg.org, and more) using a mix of Wagtail, Dokuwiki and Wordpress with the latter being phased out in favour of Wagtail.
- Wikis (wiki.ietf.org, authors.ietf.org and others) using wiki.js with custom plugins.
- Groupchat. Using Zulip.
- Mail processing is described in its own section below.
- Meetings technology is described in its own section below.
- Ephemeral notes (notes.ietf.org) using Hedgedoc
- Yangcatalog (yangcatalog.org) a custom website for a specific technology.
- We also make a considerable number of files/documents available via direct HTTP and RSYNC.

---

[2] https://www.cloudflare.com/en-gb/galileo/

All of our applications are open source and available at our main tools repository
https://github.com/ietf-tools/

Additionally, staff, contractors and leadership use a range of cloud applications, including GSuite, Slack, Salesforce, Trello, Intacct. These have varying degrees of integration with our custom services.

Mail Processing

The IETF operates hundreds of mailing lists processing tens/hundreds of thousand of messages a day and has a number of components in its mail processing chain:

- Mailman 2 is used for list management with an upgrade to Mailman 3 expected during 2023.
- All messages pass through our custom mail handler, Postconfirm which manages a global whitelist including challenge/response for new users, and rewrites addresses to avoid DMARC issues.
- The mail server itself is Postfix.
- The backend message data is stored on disk in maildir format.
- IETF participants can also access the mail archive through an IMAP server, which uses a donated licence from Isode.
- The ordinary mailman message archive has been replaced by our own archiving tool, Mailarchive, accessed at https://mailarchive.ietf.org

Meetings Technology

Our meetings technology will not form part of any RFP but is summarised here for completeness.

- Meeting registration is through the ARO system, developed and maintained by our Secretariat AMS.
- Meeting remote participation uses Meetecho, managed by Meetecho directly.
- Internet connectivity at our meetings is delivered by our NOC, a group of volunteers and contractors, using equipment donated by Cisco and Juniper.

Monitoring

Service monitoring is currently provided by a proprietary mechanism developed by our current provider. Implementing the new strategy will require this to be replaced entirely.

We use ScoutAPM for some performance monitoring of our Datatracker application and we are actively exploring similar products.

Developers

The IETF has a core set of staff developers and long-term development contractors who write the bulk of our software. Our infrastructure includes a substantial set of development/testing environments and automated build chains.

We also actively support members of the community who wish to contribute to existing projects or their own related work by building and maintaining development environments (as downloadable images), live test environments and test datasets.

Specialist Providers

The IETF is supported by a number of specialist providers in the following areas and is not expecting to review this setup: UI/UX research and design, Database management, Security auditing.

Key Information Sought

1. **Structure of the upcoming RFPs**. We are planning to issue two RFPs, one for mail processing and one for everything else. Our rationale for a separate mail processing RFP is that our mail processing requirements are quite specialised and bundling it into a single RFP would either limit those who could bid, or lead to bidders overstating their expertise in mail processing.

   a. If you consider that we could take a different approach to issuing two RFPs, then please let us know.

   b. Are there any other areas where it would make sense for us to split off a part of our infrastructure provision into a separate RFP?

2. **Separation of responsibilities**. Our in-house development team is small and highly capable. They are slowly moving into devops though with no specific target behind that. We wish to support this move while avoiding the need for them to be on call to support infrastructure providers.

   a. Is a goal of not having our team on call achievable?

   b. Do you have any recommendations for how we separate the responsibilities between our team and our service provider(s) to achieve the goal of our team not being on call?

   c. In order to achieve our goal of our team not being on call, do we need to utilise a specific deployment strategy?

    d. Do you have any other advice on how to best structure our RFPs to avoid our team being on call?

3. **Service level expectations**. Some in our community expect zero downtime from our infrastructure, but we assume this will not be possible until we have fully containerised our apps and re-architected them to enable horizontal scaling.

    a. Can we expect zero downtime from our infrastructure given that we have monolithic applications that do not horizontally scale?

    b. How should we structure our service level expectations to achieve the minimal downtime, given our current architectural constraint of monolithic applications that do not horizontally scale?

4. **Anything else**. Given that we have not tendered our infrastructure services for some years, and when we last did both our infrastructure and our requirements were quite different, is there any other advice that you would give us on our upcoming RFP(s)?

Thank you for your time.

Jay Daley
IETF Executive Director
IETF Administration LLC