Sysadmin Maintenance 2022

Request for Proposals

Summary

The KernelCI project relies on a number of web services which need constant maintenance. These include databases, automation tools and web dashboards for several instances. Some are hosted on dedicated virtual machines (VMs), others in the cloud. This Request for Proposals seeks to extend the current team with dedicated sysadmins to ensure the maintenance of these services is being carried out to guarantee a good quality of service. Additionally, some improvements can be made to reduce the maintenance burden.

Last edited	2022-06-22
Version	3
Authors	Guillaume Tucker

The KernelCI Project

KernelCI was created several years ago by Linux kernel maintainers in order to provide an automated test system for the upstream kernel. It then became a Linux Foundation project in 2019 and the founding members are still part of the project. More information can be found about it on the kernelci.org website. An annual budget is collected from the members' contributions which makes it possible to fund work packages, sysadmin maintenance being the first of this kind.

Services

The KernelCI code base provides a full system for automating kernel builds and tests. A typical instance will include the following services:

- Mongo DB database
- Backend API
- Web dashboard
- Kubernetes clusters for running kernel builds and tests
- Static VMs to run such services hosted in Microsoft Azure
- Jenkins to coordinate the automated testing pipeline
- Docker registry (currently Docker Hub)

There are currently three instances which all include the services listed above:

- linux.kernelci.org: main production instance for upstream kernel coverage
- staging.kernelci.org: development instance with changes not yet in production
- chromeos.kernelci.org: instance dedicated to running Chrome OS tests

Note: A number of test farms with hardware platforms are connected to these instances in a distributed fashion. However, their maintenance is not part of the KernelCI project's role as they belong to separate organisations and individuals.

In addition to these, there is also a KCIDB service which uses BigQuery as a database engine with an API hosted in GCE. KCIDB's purpose is to collect all the test results from KernelCI "native" instances as well as other automated systems testing the Linux kernel.

Typical Maintenance

The following maintenance tasks are typically carried out on a daily basis:

- Restart VMs or services that stop and don't recover by themselves
- Increase storage or discard old files when capacity reaches some limit
- Cancel or resubmit Jenkins jobs that fail to complete
- Investigate jobs that don't complete in Kubernetes clusters
- Monitor service costs (GCE…) to prevent exceeding the allocated budget

Occasionally, some changes need to be made to the infrastructure. Here are a few examples:

- Adding a new repository to Docker Hub

- Opening a network port on a VM to run a new service
- Installing a new service on a VM, typically deployed with Ansible
- Configuring and monitoring Kubernetes clusters
- Setting up Cloud resources (database etc.)
- Adding / removing SSH keys for users with access to the static VMs

The tasks listed above all fall within the scope of this work package and should be taken into account when responding to this RFP.

Potential Improvements

A number of things could be done to ease maintenance and improve the quality of service. Here are just a few ideas. These are not required to be part of the work package but if possible would constitute a very valuable addition. A proposal may include such improvements as separate items in a quotation.

Monitoring

A monitoring system is currently in place using xymon to send email notifications when a service or VM stops working correctly. However, it can be improved in many ways: it doesn't cover all the running VMs, there is no dashboard and the email notifications aren't always very clear.

Cloud services

The KernelCI project already has access to MS Azure and GCE clouds. Currently, Azure is used to provide static VMs to run services and Kubernetes clusters to build kernels. There are two separate GCE accounts, one for Kubernetes builders and one for KCIDB with BigQuery and its API. We could be using cloud services more, for example:

- Host Mongo DB using Atlas rather than running it directly on a VM
- Use a Docker registry in Azure to avoid the limitations of a free Docker Hub account (or upgrade our Docker Hub account)
- Gather monitoring events from such clouds into a single dashboard and notification system, which is related to the previous topic about improving our monitoring tools

Continuous deployment

At the moment, the linux.kernelci.org production instance is updated once a week by pausing all services, flushing any current pipeline jobs and ensuring services come back online correctly after the update. While we could continue to have weekly releases as this is fine from a development perspective, it would be very useful to have a more seamless production update mechanism which doesn't require a full stop. This would increase the quality of service by removing a delay in producing test results caused by the current update procedure. It is a slightly more complex issue to solve as it will most likely involve changes in how the KernelCI configuration is made (for example, Docker images would need to be tagged etc.).

Vendor Considerations

Proposed Timeline

The deadline for responding to this RFP is six weeks after it has been made public. Then the KernelCI Advisory Board of Members will vote and respond within two weeks, so a decision should have been reached within eight weeks. By starting this timeline on 27th June 2022, the deadline to respond is 8th August with a board vote on the 24th August. Exact dates might be subject to change in case of a major practical issue or unavailability of voting members.

The maintenance work package should last six months, so typically from 1st September 2022 to 28th January 2023. Alternative dates may be used in proposals and any improvements should be completed within their own defined timeframe.

Vendor Selection Criteria

Each proposal will be examined by the KernelCI Technical Steering Committee and the Advisory Board members. There is a fair amount of flexibility in this RFP which means there could be significant differences between proposals. Therefore, having a clear and unambiguous description of what is being proposed is a key factor in the decision making process. We need to be able to gain enough confidence that a proposal is a viable one before it can be chosen.

Particular points of interest are past experience in maintaining web services, managing cloud technology and some familiarity with open source software in general. The KernelCI VMs all run major Linux distributions such as Debian and Ubuntu.

Budget

We're expecting quotations for this work package to range between 15,000 and 30,000 USD depending on the contents of the proposal. Longer sysadmin time available and extra improvements can justify a higher price.

Payments would be made from the Linux Foundation using the project's own budget.

Contact information

Please send responses to this RFP via email to the KernelCI governing board:
kernelci-members@groups.io