


National Aeronautics and
Space Administration

Headquarters

Washington, DC 20546-0001



NRW 2210-35
NPR 2210.1C

NOV 17 2010

Reply to Attn of:

Office of the Chief Information Officer

TO: Office of Internal Controls and Management Systems
THRU: Office of the Chief Technologist

FROM: Office of the Chief Information Officer

SUBJECT: NPR 2210.1C Requirement Waiver in Support of Open Source Software Development

INTRODUCTION

Pursuant to NPR 1400.1D § 4.2, the Office of the Chief Information Officer (OCIO) requests a waiver from NPR 2210.1C¹ to the extent the policy requires software be completed before it may be evaluated for release. In order to support open source development of the NASA Nebula cloud computing software stack, the OCIO requests that the Nebula software's intended functionality and architecture be considered conceptually and, if the risks are deemed acceptable, that Nebula be permitted to develop the software in the open by storing its source code on publicly accessible source repositories. This will accelerate development and create the possibility of leveraging community expertise to produce higher quality software.

BACKGROUND

NASA Nebula is a project under the OCIO created to evaluate and adapt cloud computing techniques and service models for use within NASA. Cloud computing is a model for enabling convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction.² This approach to computing holds the promise of more efficiently processing and distributing NASA data in furtherance of NASA's objective to expand human knowledge of the Earth and of phenomena in the atmosphere and space.³

In reviewing the state of cloud computing in 2008, it was determined that existing solutions would not fully satisfy NASA's performance or security requirements. As such, and in

¹ This waiver relates solely to NPR 2210.1C and does not implicate other agency policy or directives.

² Peter Mell and Tim Grance, *The NIST Definition of Cloud Computing 1* (2009).

³ The National Aeronautics and Space Act, Pub. L. No. 85-568, § 102(d)(1), 72 Stat. 426 (1958) (codified as amended at 42 U.S.C. § 2451(d)(1) (2006)).

furtherance of the project's efforts to build expertise in the field, a software development effort was started to create a private cloud computing platform utilizing free and open source software. A corollary of this architectural direction has been the requirement to release developed software under free and open source licenses and to engage the open source software community in the construction of the system so as to more effectively and expeditiously advance the state and capabilities of cloud computing.

JUSTIFICATION FOR WAIVER

Under certain circumstances, NPR 2210.1C permits the release of NASA software under open source licenses.⁴ However, it has been interpreted to require that software be fully completed before it can be released. This approach has been characterized as the Cathedral model wherein source code is made available with each software release, but code developed between releases is restricted to an isolated group of developers.⁵ This approach greatly increases the cost of software development by slowing its progress and by rendering the benefits of community expertise unavailable.

Conversely, in the open source development model, also known as the Bazaar model, code is developed over the Internet in view of the public.⁶ The open source software community understands that making source code publicly available on source repositories accessible only through freely available developer tools does not constitute an official release of software, but merely publication of a project's development efforts and invitation for participation. Developers accessing code made available in this manner expect that the code is incomplete, will have flaws, and may not function at all. The willingness to expose such code to community review and development is considered a positive trait of burgeoning software projects.

The benefits of developing software in view of the public include:

- **Massively Accelerated Development.** Publicly accessible source code may be developed concurrently by any number of individuals or entities. In practice, parties participate in open source software development in furtherance of their self-interests. This has the effect of accelerating development and broadening the feature set.
- **Access to Diverse Expertise.** Every potential developer brings a unique skill set and body of experience to a project. Due to the fact that public software development has the potential to leverage a global developer base, the result is access to expertise

⁴ Generally, NASA software is released under the NASA Open Source Software Agreement. Release of NASA Software, N.P.R. 2210.1C § 3.2.2.1 (Aug. 11, 2010). However, Nebula has received dispensation to release its software under the Apache 2.0 license in order to align with community software development efforts in the area of cloud computing. Memorandum from Agency Counsel for Intellectual Prop. to Nebula Program Executive (June 28, 2009).

⁵ Eric Steven Raymond, *The Cathedral and the Bazaar* (2000), available at <http://www.catb.org/~esr/writings/cathedral-bazaar/cathedral-bazaar>.

⁶ *Id.*

covering the entire range of relevant disciplines. Amassing such a group in the Cathedral model would very likely be cost prohibitive.

- **Peer Review.** Also stemming from the broad developer base occasioned by open source development is Linus' Law, which states: "Given enough eyeballs, all bugs are shallow."⁷ In other words, "Given a large enough beta-tester and co-developer base, almost every problem will be characterized quickly and the fix obvious to someone."⁸ This has the effect of producing higher quality software at lower cost.

In addition, open source development has the added benefit of materially advancing the President's Open Government Directive which requires that federal agencies embrace, and model policy around, three core principles:

1. **Transparency.** Transparency promotes accountability and provides information for citizens about what their Government is doing. Information maintained by the Federal Government is a national asset. The Administration will take appropriate action, consistent with law and policy, to disclose information rapidly in forms that the public can readily find and use. Executive departments and agencies should harness new technologies to put information about their operations and decisions online and readily available to the public.
2. **Public Participation.** Public engagement enhances the Government's effectiveness and improves the quality of its decisions. Knowledge is widely dispersed in society, and public officials benefit from having access to that dispersed knowledge. Executive departments and agencies should offer Americans increased opportunities to participate in policymaking and to provide their Government with the benefits of their collective expertise and information.
3. **Collaboration.** Collaboration actively engages Americans in the work of their Government. Executive departments and agencies should use innovative tools, methods, and systems to cooperate among themselves, across all levels of Government, and with nonprofit organizations, businesses, and individuals in the private sector.

Specifically:

1. **Open Source Software Development is Transparent.** Open Source software licenses permit free distribution and require a program's source code be made available to all. Developing open source software publicly ensures that the purpose, functionality, and structure of software written by or for NASA is made manifestly apparent to any interested party. This inherent transparency furthers government accountability by making decision making, both at the abstract policy level through the purpose of software and at the specific code level through the structure of software, readily apparent as the software is engineered. The public is then afforded the opportunity to submit informed feedback at points in the development life cycle where such feedback can have a meaningful affect.

⁷ *Id.*

⁸ *Id.*

2. **The Public is Empowered to Participate in Open Source Software Development.** To qualify as open source, works must be amenable to modification under the terms of the applicable license agreement. This ensures that the public may directly participate in the development of government software when said software is developed in the open by making community alterations or contributions available to the Government for inclusion into relevant software projects. In this way, public knowledge is leveraged to make the best software engineering decisions available, resulting in more efficient and secure software infrastructure.
3. **Open Source Software Development Directly Fosters Collaboration.** By definition, open source software development utilizing the Bazaar model promotes an environment where public and private entities can both easily and expeditiously work together on solving problems of common interest. This is made possible by direct, multi-party discussion and code contributions focused around software projects licensed under open source software agreements.

OPEN SOURCE DEVELOPMENT PROPOSAL

The OCIO requests a waiver from NPR 2210.1C to the extent that it requires software be completed before it can be evaluated for release. In its place, OCIO requests approval to operate under the following process until subsequent policy is prepared to address the problem of open source software development or until the Nebula project is completed, but not to exceed August 11, 2015.

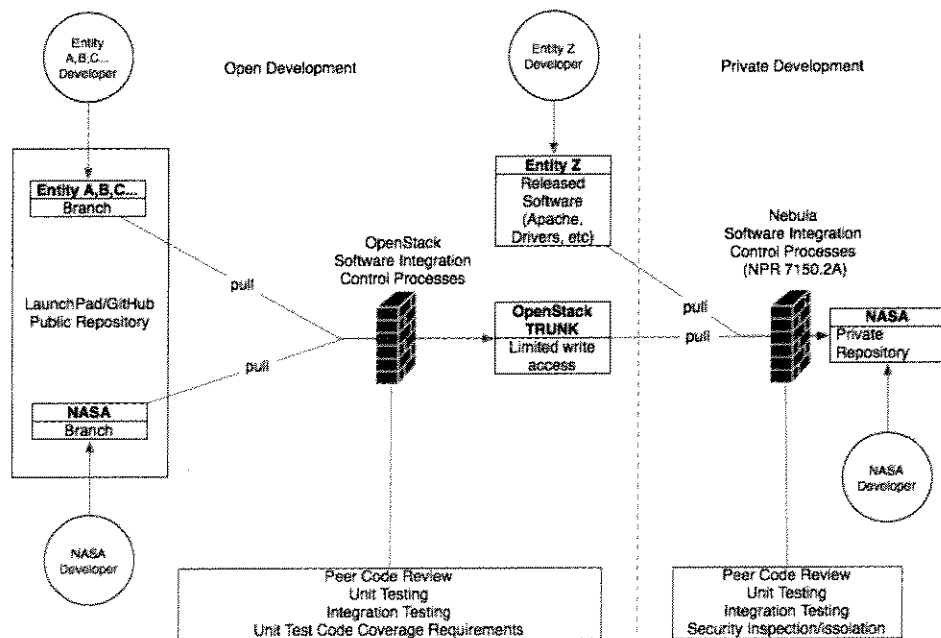


Fig. 1. Open Source Development.

At its core, open source software development rests on the ability to develop code in the public. Figure 1 describes a public-private source repository arrangement wherein generalized code is made public. Any deviations from the public code unique to NASA are stored in a private code repository. Code is always pulled from public to private – technical controls prohibit accident publication of NASA proprietary code. Each entity in the depicted system works on its own complete duplicate of the underlying codebase, making any alternations thereto publicly available for consumption by any other entity. This is always a publish-consume process. Entities may suggest consumption of their public code, but they do not push, or contribute, in the traditional sense.

As such, any entity is empowered to pull code from the NASA public repository and compatibly-licensed code may be pulled from any other entity's public repository into NASA's public repository if NASA desires the functionality enabled by such external code. Finally, Figure 1 indicates that code which is considered mature is feature frozen, assigned a version number, put through the NPR 2210.1C release process, and offered for public consumption as open source software packaged for installation by end users.

To realize this approach to software publication and development, Nebula will submit a request to the Software Release Authority describing the software's intended purpose, functionality, features, and an approximation of how the software will be implemented. The critical distinction between this approach and the existing process for software release is that the software in question need not be written. Instead, this process reviews the software's concept based on the fullest description available to determine potential liability and the applicability of export control laws. Specifically, the proposed process consists of:

1. **Umbrella Description of Software Sought to be Developed.** Nebula will submit a detailed description of the software's purpose, features, components, structure, functionality, and an approximation of how the software will be implemented at periodic intervals corresponding to significant new development activities. The software will be described as fully as possible before the software is actually written, including all information normally required in the NPR 2210.1C process such as the list of innovators, list of innovations, licensing, and related information. This description then forms the outward boundary of what Nebula is permitted to develop publicly. That is, what code may be published to public source repositories as the code is written.
2. **Legal, Export Control, IT Security, and Other Appropriate Review.** Based on the submitted software description, certain reviews of the software's concept are conducted to determine its liability profile and the applicability of export control laws. The software release authority coordinates these reviews and makes a determination as to the ability of the applying program or project to develop the software publicly within the confines of the software description.
3. **Open Source Software Development.** Once approved, the Nebula will publish source code developed within the confines of the umbrella software description to an approved public software repository as it is completed.
4. **NPR 7150.2A Compliance.** All software created or acquired by or for Nebula will meet the requirements of NPR 7150.2A. Specifically the item identified in Figure 1

by "Nebula Software Integration Control Processes (NPR 7150.2A)" will include processes to assure that the software to be integrated with Nebula meets appropriate requirements for performance, security, and integrity. These assurance methods and processes are being evaluated and developed in collaboration with the Ames Office of the Chief Engineer.

Importantly, the review described in second number item, above, is the same as that applied to completed software. All potential external software packages, as characterized in the umbrella description, are reviewed for license compatibility and acceptability; export control conducts reviews of the software and external packages; and IT security ensures that the software is developed securely and that its publication will not expose NASA assets to undue security risk. The level of reporting detail here contemplated is unchanged regardless of whether the software is written or not. Software is the expression of a concept, function, or mode of operation manifested in written instruction, procedures, and rules. The concept precedes the expression. As a matter of course, specific variations as to the expression are not requested or reviewed during the software release process – only the broader concept of the software's expression is considered. As such, limiting reporting to completed software is not necessary to achieve the required level of detail. Further, any risk associated with uncertainty as to software's fixation will be addressed and mitigated through periodic resubmissions corresponding to significant new development activities, self-enforcement of the approved development scope, and ongoing dialogs with relevant reviewing parties between periodic resubmissions to address deviations or other matters of concern.

The similarity of review notwithstanding, there is at least one important issue presented by open source software development: community developed features outside the approved software description. It may occur that the development community modifies Nebula's publicly available source code in a manner not contemplated by the software description. If such modifications are made publicly available under the same license terms which govern the original code, or under the terms of a license approved by NASA legal, it is requested that Nebula have the ability to accept such external modifications without revising the software description. To require otherwise would stymie development by requiring that any community submitted alternations to the project outside the scope of the software description require revision of the software description. In other words, acceptance of community developed features into Nebula's codebase would preclude republication of the augmented codebase as it would contain features or functionality outside of the software description. Given the rate at which open source development can occur, the paperwork and processing required to modify the umbrella software description after each desirable community deviation would likely be debilitating. Rather, such community deviations would be captured and characterized either in a subsequent umbrella description submission or in the NPR 2210.1C software release request process described below.

Due to the sensitivity of using external code obtained under varying license agreements, the Nebula team will endeavor to identify the external software contemplated for use in open development in the umbrella software description. Only external software packages and related code described in this manner and authorized for use in open development will be employed. If new, not previously contemplated external packages are deemed necessary for

development to continue, addenda will be submitted for approval by the relevant reviewing bodies before such external packages are used. Efforts will be made to ensure that the Nebula development team is fully informed as to the scope of the authorization for open development. This awareness, coupled with closely monitored and electronically tracked software development, should ensure adherence to the scope of the open development authorization. Additional controls, whether developed in-house or suggested by the Software Release Authority in consultation with relevant reviewing bodies, will be implemented to ensure conformity with the scope of the open development authorization.


RELEASE OF BAZAAR-DEVELOPED OPEN SOURCE SOFTWARE

Recognizing that software is subject to change between the time it is approved for open source development and when it reaches a level of maturity conducive to a proper release to users (as opposed to developers), this proposal suggests that the regular NPR 2210.1C release process be pursued when Nebula wishes to issue an open source software release. Key attributes of a release, as opposed to public code development, include a feature freeze, change log, version numbering, support channel, and packaging for installation. A schedule for planned public releases would be maintained as part of the Nebula master schedule with accommodation for the NPR 2210.1C review cycle. This schedule would be made continuously available to Center reviewing bodies so that adequate resource planning and time management can take place. Technical controls, such as revision control, in conjunction with developer awareness will be used to identify deviations from the approved scope to be captured in the subsequent umbrella description submission or in the NPR 2210.1C software release process.

CONCLUSION

Nebula requires approval to operate under this proposal in order to accrue the benefits inherent to open source software development. Specifically, Nebula stands at the precipice of either being able to participate in community software development or not as community advancement of components previously released under the Cathedral model proceeds apace. Development has reached the point that, if not permitted to participate in open source software development, Nebula will lose the ability to influence development efforts toward goals directly benefiting NASA, and all attendant benefits of Bazaar-style community participation will be lost. All developers will be briefed and trained on the risks associated with developing open source software, including the introduction of alternatively licensed code to ensure the success of this effort.

The OCIO will seek to formalize open source development as an option for software developers at the Agency through the promulgation of appropriate policy. This waiver is requested to grant immediate authority to undertake the requested actions.


Linda Cureton
NASA Chief Information Officer


Bobby Braun, PhD
NASA Chief Technologist


Approved