William Stein ***RRF: Cover Memorandum***

(206) 419-0925 wstein@uw.edu http://wstein.org

PI Name

William Stein

PI Department

Mathematics

Proposal Title

Salvus: Distributed Scalable Online Mathematical Software

eGC1 Application Number

eGC1 Number: A77160

Proposal Abstract

This proposal is about a web application for doing mathematical computation (based on Sage) called Salvus, which the PI is implementing from the ground up to be secure, highly available, horizontally scalable and able to robustly handle up to 1,000 simultaneous computations by March 2013 (with further scalable then being primarily a matter of available hardware resources). There will be both free and non-free accounts, with revenue from the non-free subscriptions supporting maintenance and purchase of hardware, cloud hosting, and the development of Sage. Since this commercialization effort is a major and risky new direction for the PI, he is working closely with UW's Center for Commercialization to navigate the legal and business minefields.

William Stein ***RRF: Proposed Research***

(206) 419-0925 wstein@uw.edu http://wstein.org

1 Introduction and Rationale

Most of the PI's published research work is in number theory, which is an area of pure mathematics. The PI embarked on a new and risky path at Harvard in 2005, when he started the Sage mathematical software project, which has grown substantially since he was recruited by UW in 2006. There are now around 500 Sage developers, over 10,000 downloads of Sage per month, over 55,000 unique visitors each month to the `http://sagemath.org` website, and many publications make use of Sage (see `http://www.sagemath.org/library-publications.html`). Sage is a very large free self-contained program for doing sophisticated mathematics, which you can download and run on your computer. It incorporates nearly 100 other programs, and already has interfaces to all other popular math software systems. The goal of the Sage project is to create a viable free open source alternative to the commercial systems Maple, Mathematica, Matlab and Magma. The PI has published two articles about Sage (see the citations document).

The Sage notebook, whose development was partly supported by the RRF program in 2009, is the primary user interface to Sage. It was designed mainly for use by a relatively small number of simultaneous users, but there is a public free version that anybody can quickly start using at `http://sagenb.org`. Sage is software for doing number crunching, so it is resource intensive. The `http://sagenb.org` website runs on one powerful computer that the PI purchased in **2008** and it has around 100,000 user accounts! There are often well over 50 people attempting to use `http://sagenb.org` at any given time, which can be frustrating. The core design of sagenb is not-scalable to multiple nodes, since it doesn't even use a database, and rewriting it to use one has proved extremely challenging. Also, there is currently no way for users to pay a small fee to rent more available resources.

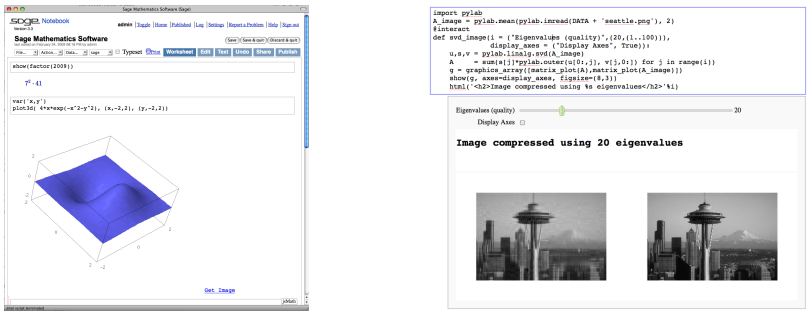

Figure 1: Screenshots Showing the Sage Notebook

(206) 419-0925 wstein@uw.edu http://wstein.org

Goal: *Create a web application (based on Sage) called Salvus, which is designed from the start to be secure, highly available, horizontally scalable and able to robustly handle up to 1,000 simultaneous computations*[1]*. Provide both free and non-free accounts, with revenue from the non-free subscriptions supporting maintenance and purchase of hardware, and the development of Sage.*

Instead of teachers and researchers having to pay to buy expensive commercial mathematical software and install it on their computers, or be frustrated with the current Sage notebook, Salvus would provide them the option to *reliably* and efficiently use Sage over the web and confidently get their students to use it. At many institutions, purchasing computer software—especially mathematical software—is a significant burden, and Sage has helped address this problem. The burden of installing, maintaining, and upgrading Sage remains, and the Sage notebook has not solved that problem because it is too slow and unreliable. This project has the potential to have a profound impact on education at all levels.

There are several commercial web-based software development environments such as `https://c9.io/` and `http://www.heroku.com/`, which have some similarities to what the PI plans to implement with Salvus. However, Salvus will be highly interactive and much more focused on mathematical algorithms, education, and research than on deploying generic web applications.

2 Objectives

The PI spent much of Summer 2012 on the foundations of Salvus, and will make a limited-functionality version available soon (see `https://salv.us`). He intends to have a version up and running with comparable functionality to the existing Sage Notebook by March 2013, and build a core subscriber base of at least 2500 users by August 2013; work during Winter 2013 is being funded by the remainder of his startup money, and during some of Summer 2013 by an NSF grant.

This RRF would allow the PI to work fulltime through late December 2013 on the next stage of the Salvus project, to address these specific goals:

1. Market an enterprise version of Salvus that customers can install on their own internal secure network. Similarly, support university and course wide site licenses for Salvus.

2. Greatly improve support for using software other than Sage through the Salvus interface, including Octave, Scilab, R (included in Sage), Matlab, Mathematica, Maple, and Magma. This will draw in users from the sciences.

3. Implement editors for at least the following document types: command line, LaTeX document, Python and Sage files, and for fast interactive 2d and 3d graphics.

[1] Far more than 1,000 users could be simultaneously connected.

(206) 419-0925 wstein@uw.edu http://wstein.org

4. Integrate with many popular cloud storage and repository hosting services, including Dropbox, Google Drive, Github, Google Code, and Bitbucket.

5. Make it so users can write code that gets included with Sage entirely using Salvus, so they can contribute to Sage without having to install Sage.

3 Procedure

3.1 Computer Hardware

A single server can serve static pages to 1,000 simultaneous users reasonably quickly. The challenge is that when users evaluate mathematical expressions, such as computation of a difficult symbolic integral or solve a large system of equations, they are potentially using substantial CPU resources and RAM. Moreover, we must keep a persistent Sage process running for the duration of their session. Salvus demands a large number of computers that have the ability to run a Sage session. For high availability, there must be machines in many physical locations.

Purchasing access to computers located at multiple data centers is critical to this proposal, given the dual goals of high availability (even when a node or whole data center goes down) and scalability (1,000 *simultaneous* computations). Use of Sage is extremely CPU intensive, and based on extensive PI experience[2], which includes dozens of hands on tutorials over 5 years using the Sage notebook, the PI estimates that appropriately implemented, n-Ghz of compute power and n-gigabytes RAM translates to a good experience for roughly $2n$ simultaneous computations. The PI is rolling out the following compute acquisition plan over the next 6 months:

1. **Padelford Hall:** 5 Dell R415's each with 64GB RAM and 16 3Ghz cores (cost of $12,500K; no hosting costs) – 480 simultaneous users.

2. **UW Tower (Colocation):** 5 Dell R415's with 64GB RAM and 16 3Ghz cores (cost of $12,500K; yearly cost of $1200) – 480 simultaneous users.

3. **Google App Engine:** the PI was awarded $60K in Google App Engine credits via a competitive application procedure. Google App Engine is a restricted application hosting environment, which can't yet run Sage; the PI is working with Google to determine how best to make use of these resources.

4. **Servedby.net cloud computing provider**[3]: 8 Linux 4R instances (about $15K/year) – 64 simultaneous users.

[2] During Summer 2012, the PI ran a 2-week "SIMUW" workshop at UW, with about 30 high school students hammering a Sage notebook server.

[3] `http://servedby.net/` is owned and run by a recent UW undergraduate.

5. **Amazon.com EC2 could hosting:** 8 m1.large spot instances (four 1 Ghz cores each). One bids on these; they are cheaper than guaranteed instances, but are not always available (about $10K/year) – 64 simultaneous users.

6. **UW backup:** 500GB of UW's Tivoli tape backup at two sites ($365/year).

The PI is focusing initially on hardware hosted at UW, since power, network bandwidth and hosting are highly subsidized by other UW activities. (Also, spare cycles on this hardware will be used for mathematics research projects while the user base for Salvus grows.)

When a node (or even a whole data center) go down, the site will continue to give users full access to their data, and they will be able to perform computations. However, users will share less compute resources, so running computations may be slower. Non-paying users will have a lower priority for resources.

If the number of paying users grows so that sufficient revenue is being generated from subscriptions, it will be easy to increase cloud computing resources to meet demands. Moreover, cloud hosting provides the option to have servers located around the world, and to dynamically grow and shrink in response to demand.

3.2 Computer Software

To support fault tolerance, the PI has designed Salvus so that it has no single (or even double!) points of failure. The underlying database is **Cassandra** (see `http://www.datastax.com/`), which is designed so that it continues to function well even if several nodes fail. Cassandra is heavily used by Netflix, Disney, and many other companies.

For internode communications security, we use the encrypted fault-tolerant virtual private network software **tinc** to connect all machines together (tinc is a P2P VPN with no single point of failure, in contrast to openVPN). For communications security between our cluster(s) and end users, we use **stunnel** and SSL, with each data center having one stunnel node and DNS pointed at all of the them. For HTTP/1.1 load balancing we use **haproxy**. We use the **Tornado** web server, which excels at serving a large number of standing connections. We serve static content using **Nginx**. The heavy lifting is done by nodes running **Sage**.

4 Time Schedule

The PI proposes to work full time for three months supported by the RRF:

1. **October, 2013:** Package and test a standalone enterprise version of Salvus. Work with lawyers at C4C to nail down licensing and copyright issues. The PI has existing longterm relationships with two large potential customers.

Improve support for R, Matlab, Magma, etc., which is of interest to such customers, and support university and course wide site licenses for Salvus.

2. **November, 2013:** Improve support for a wide range of filetypes, while integrating Salvus with popular cloud storage and hosting services.

3. **December, 2013:** Building on support for repository hosting from the previous month, make it possible for people to write and modify the code that gets included with Sage entirely using Salvus. This will indirectly contribute to the ultimate goal of improving the core capabilities of Sage.

5 Need for RRF Support

Though the PI has obtained funding for the Sage project from Microsoft, Google, the National Science Foundation, prize money, private donations, and contracts for Sage development, this has mostly been limited seed money. The PI is the main architect and programmer of Salvus, and currently has no source of funding so that he can work more on Salvus during the next academic year. Salvus subscriptions may eventually provide such funding, but this RRF would address the critical period before the subscriber base grows sufficiently.

As explained above, the goals of the current proposal are to greatly enhance Sage's level of *robust availability* and provide a sustainable UW-based revenue stream to support Sage development. If successful, this will have a huge and direct impact both on development of Sage and the community of Sage users. This project thus directly addresses the mission of the Royalty Research Fund by providing a unique opportunity to increase the PI's competitiveness for subsequent funding via individual user and site-wide subscriptions to Salvus.

The last RRF that the PI received in 2009 helped make Sage attractive to the *educational market*, which is a huge area of potential funding that Sage had not had success in yet. The PI subsequently received a CCLI Type 2 grant (see `http://utmost.aimath.org/`) from the educational part of NSF, which supported the key first month of work on Salvus during summer 2012.

5.1 RRF Scholar Application: documentation of teaching load

In a typical year, a Professor of Mathematics at University of Washington would teach a total of 3.3 "points" of courses, spread across the three quarters: Fall, Winter, and Spring. The points are determined as follows:

1. 0.95 points for Math 111 (Algebra with Applications), 112 (Application of Calculus to Business and Economics)

2. 0.85 points for Math 120 (Precalculus), 126 (Calculus with Analytic Geometry III)

3. 0.8 points for Math 170 (Mathematics for Elementary School Teachers), 124/5 (Calculus), 134/5/6 (Honors Calculus), 334/5/6 (Advanced Honors Calculus), 381 (Discrete Mathematical Modeling)

4. 0.65 for other 300 level courses (e.g., Matrix Algebra with Applications), and 171 (Mathematics for Elementary School Teachers)

5. 0.76 points for 400 (senior undergrad, e.g., Numerical Analysis) and 500 level graduate courses (e.g., Modern Algebra), 380 (Intermediate Topics in Undergraduate Mathematics), 394/5/6 (Probability).

Every year one teaches different courses. What is assigned to teach depends on preference, what is available, student interest, etc.

The PI is requesting replacement salary for 2 Calculus courses, e.g., MATH 124: *Calculus with Analytic Geometry I*, 5 credits.

William Stein ***RRF: Budget***

(206) 419-0925 wstein@uw.edu http://wstein.org

01 Salaries	
William Stein, Professor, replacement salary cost estimate, Fall 2013 (RRF Scholar)	$15,000
Subtotal	$15,000
02 Contract Personal Services	
Subtotal	0
03 Other Contractual Services	
Subtotal	0
04 Travel	
Subtotal	0
05 Supplies and Materials	
Subtotal	0
06 Equipment	
Subtotal	0
07 Retirement and Benefits	
01-10 William Stein, Professor (RRF Scholar), $\$15,000 \times 26.9\% = \$4035$	$ 4035
Subtotal	$ 4035
08 Operating Fee/Tuition	
Subtotal	0
TOTAL BUDGET	**$19035**

Justification: The replacement salary will reduce the PI's teaching load by two calculus courses, which will allow him to work full time on the proposed research for one quarter.

William Stein *RRF: Curriculum Vitae*

(206) 419-0925 wstein@uw.edu http://wstein.org

Professional Preparation

Harvard University	NSF Postdoc, 2000–2004
University of California at Berkeley	Mathematics, Ph.D. 2000
Northern Arizona University	Mathematics, B.S. 1994

Appointments

- Professor of Mathematics, University of Washington, September 2010–present.
- Associate Prof., University of Washington, April 2006–September 2010.
- Associate Prof., UC San Diego, July 2005–April 2006.
- Benjamin Peirce Assistant Prof., Harvard University, July 2001–May 2005.
- Clay Mathematics Institute Liftoff Fellow, Summer 2000.

Synergistic Activities

- **Research Tools:** Founder and director of Sage (`http://sagemath.org`), which is a large free open source software project (well over 100,000 lines of the Sage code was written by the PI). Author of much modular forms, modular symbols, and modular abelian varieties code included in Magma.
- **Outreach:** SIMUW 2006, 2007, 2008, 2012 math camps; Canada/USA MathCamp mentor (2002); Math Circles talks in Boston; 2011 REU on elliptic curves; involved many undergraduates in work on the Sage software.

Graduate and Postdoctoral Advisors

- **NSF Postdoctoral advisor:** Barry Mazur, Harvard University.
- **Ph.D. advisor:** Hendrik Lenstra, University of Leiden, Netherlands.

Thesis Students:

- **Thesis Students** 4 Ph.D. students at UW: Robert Bradshaw's 2010 Ph.D. on *Provable Computation of Motivic L-functions*; Robert Miller's 2010 Ph.D. on *Computational Verification of the Birch and Swinnerton-Dyer Conjecture*; currently advising Alyson Dienes's Ph.D. thesis on *Modular degrees of elliptic curves over totally real fields* and Simon Spicer's Ph.D. thesis work. Advised eight undergraduate senior theses at Harvard and three at UW.

William Stein RRF: Other Research Support

(206) 419-0925 wstein@uw.edu http://wstein.org

This is a list of the sources, dates, amounts, titles and relevance of research awards and contracts for which Stein was the PI during the past three years.

- (pending) NSF, 2013, \$159,972, *FRG: Collaborative Research: L-functions, modular forms, and related objects*; fundamental research in number theory unrelated to current proposal.
- NSF, 07/01/12 - 06/30/15, \$224,557, *Explicit Approaches to Elliptic Curves, Modular Forms, and Modular Abelian Varieties*; fundamental research in number theory unrelated to current proposal.
- NSF, 06/01/12 - 05/31/15, \$97,114, *Collaborative Research: Sage-combinat: Developing and Sharing Open Source Software for Algebraic Combinatorics*; fundamental research in combinatorics unrelated to current proposal.
- NSF, 09/01/10–08/31/13, \$62,706, *Collaborative Research (AIM): UTMOST: Undergraduate Teaching in Mathematics with Open Software and Textbooks*; provides summer support during which I did work on Salvus, and provides valuable opportunities to interact with potential customers.
- NSF, 09/15/10–08/31/13, \$140,000, *Sage: Unifying Mathematical Software for Scientists, Engineers and Mathematicians*; supports workshops on Sage, which is one of the software components that is relevant to the current proposal.
- US Department of Defense, 09/01/10 - 09/20/13, \$209,664, *SAGE: Mathematical Software Development*; supports bug fixing workshops on Sage, but is otherwise unrelated to the current proposal.
- NSF, 2008–2012, \$1,219,078, *FRG: L-functions and Modular Forms*; fundamental number theory research unrelated to the current proposal.
- NSF, 2007–2010, \$144,543, *The Sage Postdoc* (a 3-year postdoc at UW); supported some development work on Sage.

William Stein *RRF: Literature Citations*

(206) 419-0925 wstein@uw.edu http://wstein.org

This a selection of the PI's recent publications that are relevant to this proposal. They are on Sage or use Sage extensively for applications to education and research in number theory.

1. *A Database Of Elliptic Curves Over* $\mathbf{Q}(\sqrt{5})$ *– First Report (16 pages)*, with J. Bober, A. Deines, A. Klages-Mundt, B. LeVeque, R. A. Ohana, A. Rabindranath, P. Sharaba, 2012, to appear in ANTS X. (Used Sage extensively and involved many students.) `http://wstein.org/papers/sqrt5/`
2. *Sage: Creating a Viable Free Open Source Alternative to Magma, Maple, Mathematica, and MATLAB* (9 pages), Foundations of Computational Mathematics 2011 proceedings (the PI was a plenary speaker). `http://wstein.org/papers/focm11/`
3. *The Sage Project: Unifying Free Mathematical Software to Create a Viable Alternative to Magma, Maple, Mathematica and Matlab* (16 pages), with B. Erocal, 2010, for the PI's plenary talk at the 2010 International Congress of Mathematical Software in Japan. `http://wstein.org/papers/icms/icms_2010.pdf`
4. *Fast Computation of Hermite Normal Forms of Random Integer Matrices* (9 pages), 2010, J. Number Theory. (Paper that came out of implementing core linear algebra algorithms in Sage.) `http://wstein.org/papers/hnf/`
5. *Verification of the Birch and Swinnerton-Dyer Conjecture for Specific Elliptic Curves*, with G. Grigorov, A. Jorza, S. Patrikis, and C. Patrascu (26 pages), 2009, Mathematics of Computation. (Used Sage greatly; involved many students.) `http://wstein.org/papers/bsdalg/`
6. *Elementary Number Theory: Primes, Congruences, and Secrets* (book), 2008, published by Springer-Verlag. (Extensive Sage code included.) `http://wstein.org/ent/`
7. *Explicitly Computing Modular Forms* (284 page), Graduate Studies in Mathematics (AMS), 2007, with an appendix by Paul Gunnells. (Extensive Sage code included.) `http://wstein.org/books/modform/`
8. *Opinion: Open Source Mathematical Software* (2007), with D. Joyner, Notices of the AMS, **54**, no. 10. `http://www.ams.org/notices/200710/`