

HOW TO DO CORPORATE RESEARCH

Ken Haase

v0.81 4 February 2005

This document describes how to create and manage a corporate research laboratory to ensure continued creativity and productivity while avoiding isolation and irrelevant incrementalism.

Contents

1 INTRODUCTION	1
2 HOW TO ORGANIZE RESEARCH	2
3 HOW TO HIRE RESEARCHERS	3
4 HOW TO ORGANIZE RESEARCHERS	3
4.1 Rules of Thumb	4
4.2 Research Pensions	4
5 HOW TO ENGAGE WITH THE LARGER COMPANY	5
6 HOW TO COLLABORATE OUTSIDE THE LAB	5
7 HOW TO EXPOSE AND PROTECT THE WORK	6
7.1 Intellectual Property	6
7.2 Publication and Presentation	7
7.3 Open Source Releases	7
8 HOW TO MANAGE SPACE & TIME	8
8.1 Location, Location, Location	8
8.2 Architecture	8
8.3 Rhythm	9

1 INTRODUCTION

This document describes how to create and manage a corporate research laboratory to ensure continued creativity and productivity and avoid isolation and irrelevant incrementalism.

It discusses how to organize research activity, how to hire researchers, how to engage with the larger organization, and how to expose the laboratory to the larger world. It is based on the idea of optimizing four distinct principles that we think are necessary for the long term success of a corporate research laboratory:

- **be fresh:** encourage churn and experimentation with ideas and perspectives, keeping a diverse portfolio of research activities;

- **be relevant:** align research activities with the company's broad mission though not solely with its current products or objectives;
- **be excellent:** bring together brilliant people and challenge them to do look deeply and creativity at hard problems;
- **be nice:** explicitly, by hiring and incentive, create a *culture of generosity* where researchers enable the creativity of others throughout the organization.

None of these recommendations are intended to be rigid guidelines. In general, there should be an aesthetic of minimizing structure and letting form follow function. We also offer some examples from a new laboratory we are currently designing.

2 HOW TO ORGANIZE RESEARCH

Nicholas Negroponte has said that "managing research" is almost an oxymoron. While this is somewhat true in the academic context, where the horizon and audience is unbounded, it is less true in the corporate context, where the context for the research is more particular. The challenge is how to align research with corporate goals without either stifling creativity or running the research engine down the same ruts dug in by day-to-day production.

As we think about how to manage and organize research, it is important to work within the right metaphorical frame. Managing research is not like engineering in a conventional sense, but is more like agriculture, or husbandry, or venture investment: the goal is aggregate output achieved by individual activities.

The goals of being fresh, being excellent, and being relevant imply focused activity, active reflection, and continual opportunities to identify and learn from failure. This leads to an overall organizational scheme characterized by broad themes and particular activities.

The broad themes should be clearly aligned with the corporate mission and should provide both challenge and focus. We suggest that the themes be articulated as **core questions** which are profound, provocative, and playful. This suggestion is based on the practice of Charles Vest, when he was president of MIT, of asking faculty and other researchers to identify "things we don't know" as a way of articulating and driving research agendas.

Concrete and focused **missions** are inspired by these core questions, which will slowly evolve and change with technologies, markets, and staff. Missions are research activities with leaders, teams, budgets, milestones and deliverables. Missions are reviewed and adjusted on an ongoing basis and fall into two broad categories: prototype missions and capability missions.

Prototype missions are about creating working applications or artifacts which embody partial answers to the core questions and which provide informative experiences and technological challenges. Capability missions are about robustly realizing the capabilities necessary to build the prototypes within the laboratory. Both sorts of missions may yield products and services of commercial value, but much of their value is in helping the rest of the organization expand the bounds of the possible.

An important element of this model are twice-yearly mostly internal reviews of each mission which are used to articulate results, understand failures, reevaluate resources, and reflect on and modify milestones and deliverables. These reviews, while they should be kept as unobtrusive as possible, are central to the goals of excellence, relevance, and freshness.

In addition to formal missions, researchers are encouraged to spend up to 25% of their time on pet projects which draw on the standing resources of the laboratory and may lead to new missions or even open up new core questions. All researchers have small discretionary budgets to spend and share to support these informal but quietly crucial activities.

3 HOW TO HIRE RESEARCHERS

At all levels, the laboratory should focus on hiring *brilliant generous builders* as researchers. It goes without saying that the laboratory needs to hire the best and the brightest. However, we believe it is just as important to hire people who are *naturally generous*. There are two key reasons for this: first, generous researchers will naturally collaborate with other researchers, which will be central to the laboratory's success; second, researchers will need to actively collaborate with the rest of the organization to make sure that the laboratory's work is and stays relevant to the organization.

Finally, we believe that it is important to focus on people who build things, whether they be engineers, descriptive linguists and anthropologists, or experimentalists in the physical or human sciences. There are three core reasons for focusing on building things:

- it provides a clearer basis for ongoing evaluation and review to ensure excellence and freshness;
- it helps to ensure relevance because the rest of the organization may be able to use what was built;
- it helps cross disciplinary divides; ecumenicals have a slogan: dogma divides, service unites; building things allows different disciplines to communicate, collaborate, and learn from each other in ways they otherwise would not.

Is it a mistake to hire for a personality characteristic such as generosity? We don't think so and take a lesson from Southwest Airlines, which has an explicit policy of hiring for attitude and training for skills. Their argument, which we think makes sense, is that attitude is the hardest thing of all to change. Generosity and the desire to enable others to do more is central to the need for a corporate research laboratory to have an impact on its corporation. Furthermore, the world is blessed with enough smart people that we won't come up short by focusing on the generous ones.

A final note on hiring introduces the novel idea that, in the laboratory, **everyone does research**. This rule applies to the formal research staff but also to the technical services staff and the general support staff. Everyone, in one way or another, has active involvement in research missions and make contributions which are visible and identified. This makes for diversity in the research perspective, improves coordination generally, improves motivation, and will tend to attract overqualified support staff, which will make everything work better.

4 HOW TO ORGANIZE RESEARCHERS

One of the recurring problems in non-academic research laboratories is the problem of "too many chiefs". Such laboratories hire really smart people who can readily generate more interesting ideas than they have the attention or resources to handle. The result is a laboratory with a breadth of activity where remarkably little gets done.

One of the other recurring problems in such laboratories is intellectual diffusion. As brilliant creative minds explore and invent (and compete in a hopefully friendly way), the research focus naturally diffuses. The more researchers, the more quickly that diffusion happens, undermining the quest for relevance and the possibilities for the very collaborations the laboratory is trying to create.

In order to address these two problems, we will hire researchers in three broad categories: gurus, jugglers, and hackers.

- **gurus** do research, start and lead missions, but are also particularly responsible for shaping and refining the lab's broader themes, the *core questions*;

- **jugglers** do research, start and lead multiple missions, and do much of the day-to-day management of research activity and the progress of particular missions;
- **hackers** do research focused mostly on one particular mission, while kibitzing and learning from other missions and pet projects as their skills or interests dictate.

The gurus are senior researchers who shape and refine the core questions, together with corporate management and under the umbrella of the broad corporate mission. They are practiced visionaries who set an example by staying active in research and engineering but have a special responsibility for nurturing the intellectual and social community of the laboratory.

The jugglers are managing researchers who create and manage the concrete missions which grow out of the core questions. Like everyone else, they do practical building and research, but their activity is spread among multiple missions, ensuring their progress and evolution.

The hackers do most of the work, focusing on a single mission and doing the working and learning to make it happen. In many cases, hackers will become jugglers or move into the larger organization. Gurus and jugglers are hackers at heart, but they also invest energy in the context and the envelope of the work.

From an HR perspective, there will be a conscious effort to ensure slow but steady churn among the hackers, transforming them to jugglers, sending them off to school, or encouraging them to join the mainstream engineering and production activities. In many cases, hackers will be in degree programs and working at the laboratory as interns or as part of their degree programs. This will lead to a natural degree of churn in the laboratory.

4.1 Rules of Thumb

The general structure of gurus, jugglers, and hackers is intended as a way to separate out the roles of long term direction, research management, and focused execution. But by hiring builders and minimizing meetings and other overhead, the expectation is that all the research staff spend most of their time doing research: building things, staying on top of multiple fields, figuring things out, and observing how things are used by users and developers alike. On top of this, 20% of people's time will be spent on pet projects which use the laboratory's standing resources but are not (yet) formal missions.

Staff Distribution. Gurus, jugglers, and hackers are in a population hierarchy: roughly 5 jugglers for every guru and 5 hackers for every juggler. This is not a reporting structure: hackers report to the gurus and jugglers who are their mission leaders (plural) *ex officio*, rather than as individuals. There are, however, designated mentoring relationships across the organization and both these relationships and team reviews will factor into personnel processes.

Compensation. Compensation for gurus and jugglers should be highly competitive, though it may involve significant non-cash elements. Compensation for hackers should be competitive, but we want to have activity and hire gurus and jugglers who create a strong incentive in recruitment. Compensation for hackers will be especially variable, based on skills and experience. *Importantly, compensation should not be the reason that a hacker becomes a juggler.*

Churn. We mentioned above that we want to encourage churn within the pool of hackers and we will have explicit target for churn with hopes that folks move into production, go to school, or become jugglers.

4.2 Research Pensions

Churn and turnover are realities of any organization. We recognize this, but we also want to sustain relationships with researchers after they move on and we (frankly) also want to encourage them to move on to academia or non-profit research (where they benefit the world) rather than to our competition.

We also want researchers to be encouraged to follow some very long-term (multi-decade) research leads.

To achieve both of these goals, we introduce the novel idea of a "research pension". All researchers accrue a certain amount into a "research pension fund" for every year they work at the laboratory. When they leave for academia or non-profit research, providing other conditions are met (corporate access to research results among them), they will get a research grant from the accumulated research pension fund to continue parts of the work they may have begun at the laboratory.

Research pensions are a way to sustain valuable relationships and to encourage researchers to consider decades-long research projects.

5 HOW TO ENGAGE WITH THE LARGER COMPANY

One of the biggest problems for corporate research laboratories is isolation from their corporate context. Because they are chartered to "take the long view" and "think out of the box," there is a natural tendency for attention and action to drift from the business of the company and for the company to view the laboratory as pampered outsiders. This isolation makes it difficult for the company to benefit from the laboratory and, when times get tough and finances get tight, often leads to closure or gutting of the laboratory.

In this section, we describe four ways that a research laboratory can stay engaged with the company as a whole, avoiding the all-too-common problems of isolation. There are others, but these four make a good starting point:

- **internal sabbaticals:** staff on the engineering and production side should be actively encouraged to take internal sabbaticals at the research laboratory, where they will partake in the culture of the laboratory and actively participate in and contribute to research missions;
- **internal consulting:** research staff will actively consult on problems and issues on the engineering and production side of the company, bringing their skills, expertise, and the laboratory's emerging technologies to present business problems; this internal consulting will not be compensated, but it will be accounted for in some manner;
- **webcast colloquia:** the research lab will naturally host regular internal and external colloquia on a range of topics; these will be webcast throughout the organization and made available online archivally. Online forums after each colloquium will be open to the entire organization (or at least virtual attendees);
- **technology workshops:** as the laboratory develops prototypes and capabilities, researchers will go into the larger corporation to offer hands-on workshops around the prototypes and technologies.

6 HOW TO COLLABORATE OUTSIDE THE LAB

A laboratory cannot hire all the expertise and creativity it needs. This is partially due to practical concerns such as resources and the willingness of people to move or relocate. But it is also partially due to the fact that we want some expertise selectively. For example, we might want input from brilliant people in theory of computation, for instance, but we may not want to do lots of basic research in that area. Finally, our general hiring policy (brilliant generous builders) may rule out some people who we really want to have help us. The solution to all of these questions is outside collaborations of various sorts.

In thinking about these outside collaborations, we include both academic collaborations and collaborations with the large and growing number of non-profit technology-oriented NGOs. Thus, we will look at projects with partners like EFF, the United Nations, Computer Clubhouses, etc.

There are two broad categories here: bringing people in to the lab and funding work outside the lab. In the first area, we may bring people in for fixed-term visits or sabbaticals, where staff from other places (at all levels from junior to senior to executive) come to the lab to spend months or a year engaged with the work, covered by NDAs as appropriate.

Also in this area, we will bring needed experts into the lab as consulting gurus, jugglers, or hackers. This consulting should typically be limited to the cases of special expertise or to people whom (for whatever reason) cannot be brought on as research staff. In addition, the lab should actively encourage consultancy which is tied to other kinds of research relationships, such as funded outside research.

The lab should also fund some outside research for some of the same reasons is supporting visitors. It brings important expertise, intellectual diversity, and a broadened community. However, the way we engage should be as focused on people as it is on institutions. In particular, we should prefer to support research where there is a parallel consultancy relationship (and expect consulting fees to be discounted as a result) and we should also prefer to support research through mechanisms such as student fellowships rather than directed research contracts.

Much or most of our external funding should be directed towards external efforts which have some public good as part of their mission, such as the creation of open source software or databases. This encourages traditional academic freedom and publication, strengthens the company's reputation and brand, and lessens the impact if results or findings "leak" into other companies.

Finally, if we can manage to have our external funding be charitable giving (with a side-benefit to the company and others), that would be great.

7 HOW TO EXPOSE AND PROTECT THE WORK

This section discusses how to expose the laboratory to the outside world, given the desire to give a particular advantage to the corporation itself while both doing good in the world and helping burnish the corporate brand. We discuss exposure in three broad categories: handling intellectual property, publication and presentation, and open source releases.

7.1 Intellectual Property

The laboratory will naturally produce a range of valuable intellectual property, ranging from patentable inventions to copyrightable code to trade secret knowledge to proprietary databases. One of the lab's top priorities will be to make sure that intellectual property is protected in a timely and appropriate way. The laboratory will start with three particular mechanisms to ensure this:

- a fast, low-overhead, and clearly articulated process for getting intellectual property classified;
- general education about intellectual property, including patent and copyright issues, as well as the merits or complications to publication or open source release;
- explicit consideration of protection issues in the biannual mission reviews, identifying new or expected intellectual property which will need to be considered.

Intellectual property issues, particularly around patents and proprietary code, are often hot buttons in academic and hacking circles. The laboratory needs to have a balanced and visible consideration of these issues in all their complexity while also having clear policies in place and documented processes for handling gray areas.

7.2 Publication and Presentation

Research laboratories vary widely in their policies towards scholarly publication and public presentation, ranging from forbidding most of it to encouraging lots of it. We think that a middle ground makes sense, allowing publication which doesn't undermine the company's intellectual property, encouraging publication which can promote informative discourse, but avoiding the idea that publication is an end in itself.

Submission for publication and publication itself will require approval to make sure that any exposed technologies are adequately protected. For the role of publication in personnel review, we follow three simple guidelines:

- publications can count positively in review processes
- non-publication will not be counted negatively in review processes
- excessive publication will be looked on with suspicion

These guidelines reflect the fact that academic publication is a complex system which is heavily gamed, and we don't want to waste time and energy playing that particular game while still wanting to encourage researchers to establish themselves and their reputations.

Regarding presentations, whether connected to publication or not, it will be permitted and sometimes encouraged in much the same way as publication. A significant special case is that of presentation in artistic contexts, which may often emerge from creative people playing with new technologies. This will be actively encouraged, especially in the form of pet projects which use our new technologies in new ways.

7.3 Open Source Releases

The foundation of the Internet consists of open source software created by brilliant individuals and dedicated to the common good. We recognize our debt to this legacy and want to repay it by our own contributions, subject to the interests of the company. There are four important reasons for releasing open source software:

- it does good in the world by enabling the creativity of others;
- it exposes technology to a world of creative developers and debuggers;
- it can establish defacto standards in areas where the company is working;
- it commodotizes product categories to strategically move competition towards our corporate strengths.

Against those four reasons lies the fact that the company's present and future competitive advantage will partially lie in its software and what the software allows the company to uniquely do. For software emerging directly from capability and prototype missions, we will think about open source release in parallel with other forms of protection. Often, a mission will be started with deliverables and milestones that may explicitly include open source releases. In other cases, we will clearly identify technologies which we think it is important to keep proprietary. We will try to make such decisions as early as possible in order to ensure internal transparency and provide for future releases.

We will also devote energy to being creative in the ways that we use various forms of public release. For example, we may consider different kinds of open source licenses for different technologies and, in some cases, draft new licenses for particular kinds of technologies to ensure certain effects and returns to the laboratory and the company.

8 HOW TO MANAGE SPACE & TIME

Location, architecture, and rhythm are all important determinants of the atmosphere, style, and success of a research laboratory. This section looks at these issues and some rules of thumb for dealing with them.

8.1 Location, Location, Location

Where should a research lab be located? There is a school of thought which suggests that research labs should be isolated far away from the daily buzz of human activity, allowing their researchers to think deep thoughts.

This is a bad idea for a corporate research lab, particularly one that is trying both to remain relevant to its host and to attract and inspire young researchers. Instead, the location of a research laboratory should be close to the places where they live their lives. We want to have researchers who are passionate enough that they will dedicate a weekly 60 hours to their work out of love, liberally intermingled with the many other things they enjoy. Distance makes it hard to do that and so we can identify a few guidelines for location. A research laboratory location should:

- be urban or semi-urban
- be near a diversity of dining, pubs, and other activities
- be near some sort of natural escapes (parks, rivers, ocean)
- have affordable residential options within bicycling distance
- be convenient to public transportation
- have guest accommodation very nearby

The idea here is to have a place that is fun to be while providing the necessary escapes for the full and nuanced life which supports happiness and creativity.

8.2 Architecture

I am not an architect, but I have been involved in a lot of discussions with architects about laboratory space. The following suggestions are based on those discussions and the experiences which preceded or followed them. In general, the office layout should be an open plan of desks with attractive computers, probably clustered but without partitions. This general layout should be punctuated by:

- meeting spaces, some with walls and some acoustically isolated without walls
- dedicated conference rooms for the directors and each senior researcher (their formal desks will be in the open plan)
- work bubbles, with walls, doors, loaded workstations, whiteboards, table space, and the ability to comfortably seat three and cosily seat four.
- a single "playroom space" with both a large open area and small tables
- a quiet cafe with wireless access; this might best be a public facility run by some sympathetic entrepreneur but located on or adjacent to laboratory premises.

The idea here is to have a broad range of interaction with ample opportunities for focused activity. One of my biggest early mistakes in on laboratory was accepting the distant office with a door that my predecessor had planned to move into before I arrived.

8.3 Rhythm

The rhythm of academic research labs is typically set by the academic calendar and the schedule of thesis and conference deadlines. The rhythm of a corporate research lab is often more disjointed, which can impede progress.

In our research lab, a natural rhythm will come out of the biannual mission reviews, which will be arranged so that related missions come together while disparate missions are staggered. This will intertwine with an annual retreat and overall research review, which will include all of the researchers and visitors from around the company.