



IBM Center for
The Business of Government

Collaborating Across
Boundaries Series

Designing Open Projects

Lessons From Internet Pioneers



David Witzel
Fellow, The EdgeLab

Designing Open Projects: Lessons From Internet Pioneers

David Witzel
Fellow, The EdgeLab

Table of Contents

Foreword	4
Introduction	6
What is an Open Project?	8
A Selective History of the Internet's Development	13
ARPANET	14
The Network of Networks	17
The World Wide Web	19
Web 2.0	20
The Next Wave	21
Lessons from the Internet	22
Tip One: Let Everyone Play	22
Tip Two: Play Nice	23
Tip Three: Talk About What You Are Doing While You Are Doing It	25
Tip Four: Use Multiple Channels of Communication	26
Tip Five: Give it Away	27
Tip Six: Reach for the Edges	29
Tip Seven: Make it Work, Then Make it Better	30
Tip Eight: Make it Work, Then Standardize	31
Tip Nine: Take Advantage of All Types of Organizations	32
Tip Ten: Design for Participation	34
Tip Eleven: Increase Network Impact	36
Tip Twelve: Build Platforms	38
Case Study: Open Project Design at the Office of the National Coordinator for Health Information Technology	40
Background	40
CONNECT Software	41
NwHIN Direct	42
ONC Results: A Preliminary Assessment	44
Conclusion	45
Acknowledgements	46
About the Author	47
Key Contact Information	49

Foreword

On behalf of the IBM Center for The Business of Government, we are pleased to present this report, *Designing Open Projects: Lessons from Internet Pioneers*, by David Witzel.

This report offers practical design advice to public managers and political leaders who are facing complex, dynamic public challenges involving multiple stakeholders on issues or problems where there is no clearly defined solution. In these situations, open project approaches have the potential to spark large-scale activity that could fundamentally change society.

The author, David Witzel, examines the evolution of the Internet over the past four decades, exploring how a wide range of autonomous, overlapping, and interconnected open projects initiated by government staff, techies, entrepreneurs, and students around the world resulted in one of the most profound changes in society across the globe since the dawn of the Industrial Age. How did it happen? How was it guided? Are there lessons that can be drawn to tackle other enormous societal challenges?

Based on his observations about the creation and evolution of the Internet, Witzel identifies a dozen tips for designing open projects. Witzel believes, on reflection, that these 12 practices were critical to the Internet's successful development. He then takes those design elements and shows how they are now being applied in other domains, such as the creation of the nationwide health information network involving over 800,000 doctors and hospitals.



Jonathan D. Breul

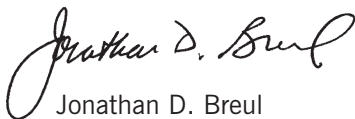


Daniel Prieto

A key insight underpinning Witzel's tips is that this is not a precise methodology to be followed. Instead, an open project approach should be viewed as a mindset. Leaders have to discern whether the challenges they are facing can best be solved using a closed or open approach, defined as follows:

- A *closed* project has a defined staff, budget, and outcome; and uses hierarchy and logic models to direct activities. It is particularly appropriate for problems with known solutions and stable environments, such as the development of a major highway project.
- An *open* project is useful to address challenges where the end may not be clear, the environment is rapidly changing, and/or the coordinating entity doesn't have the authority or resources to directly create needed change. In these *open* projects, new stakeholders can join at will, roles are often informal, resources are shared, and actions and decisions are distributed throughout the system.

Witzel's report provides guideposts on how to use an open project approach on appropriate large-scale efforts. We hope this report serves as an inspiration and practical guide to federal managers as they address the increasingly complex challenges facing our country that reach across federal agency—and often state, local, nonprofit, and private sector—boundaries.



Jonathan D. Breul
Executive Director
IBM Center for The Business of Government
jonathan.d.breul@us.ibm.com



Daniel Prieto
Vice President
Public Sector Strategy & Change
Daniel.Prieto@us.ibm.com

Introduction

The growth of the Internet is amazing. In just 40 years, it has developed from four computers connected by telephone lines into a wired and wireless service connecting over a quarter of the people on the planet and beyond. It has helped transform (and often doom) entire industries, including book, record, and DVD stores, newspapers, the Yellow Pages, and the print version of the Encyclopedia Britannica. McKinsey & Company calls the Internet “a significant and growing portion of global GDP. Indeed, if measured as a sector, Internet-related consumption and expenditure is now bigger than agriculture or energy.”¹

The vast tool set built on the Internet—from communications to analysis to data retrieval to music and video—has made us smarter, faster, more talented, and better connected. As we learn how to better use and keep improving these rich tools, we will get better at working together, managing projects, and solving problems in both government and the private sector.

But being able to use the Internet is not the whole story. Easy to overlook is the fact that the Internet itself represents a very large, distributed, collaborative development project. It has the following characteristics:

- No one owns or manages it
- It can't be turned off (at least not yet)
- Much of its infrastructure is available for reuse for free, by anyone in the world
- It is both public and private, global and local, commercial and cooperative
- It evolves, grows, breaks, and is repaired without central coordination

The Internet constitutes what can be called an open approach in two senses of the term. From a computing perspective, it is driven by public standards that support interoperability with a wide range of hardware and software products. From a systems perspective, it is flexible with permeable borders, continuously interacting with and gaining energy from its context. It was constructed via a host of often overlapping, often open projects that built upon, extended, and replaced services and created new resources and possibilities in rapid cycles of innovation and creation.

The Internet has spurred a great deal of creativity and innovation, reached considerable scale, and had a big impact on society—three things that government and nonprofit leaders strive for. Insights from the Internet's pioneers can help us create and support open projects, defined as projects that:

- Evolve, react, and innovate

1. Matthieu Péliissié du Rausas, James Manyika, Eric Hazan, Jacques Bughin, Michael Chui, and Rémi Said. “Internet matters: The Net's sweeping impact on growth, jobs, and prosperity.” Report, McKinsey Global Institute, May 2011. http://www.mckinsey.com/mgi/publications/internet_matters/index.asp

- Engage large numbers of people in learning, teaching, and creating
- Enable insights, solutions, and resources to both solve problems and help problems get solved

Internet-like open projects are now beginning to shape government and nonprofit efforts. One in particular is the Department of Health and Human Service's Nationwide Health Information Network (NwHIN), what some call a "health internet." This project is creatively applying Internet lessons to address a set of particularly thorny health and health care information issues. The NwHIN experience provides a clear example of how lessons from the development of the Internet can help government and nonprofit managers approach their own complex problems. A case study of development of the NwHIN by the Department of Health and Human Services (HHS) is presented to explore how these tips are now being used in a real-world program, and specifically how they may be applied to improve health care in the U.S.

This report will describe the idea of open projects, provide a history of the Internet's development, and offer a set of handy tips suggested by its development. These tips are intended to help government and nonprofit leaders design and implement open approaches in the projects they design and manage.

What is an Open Project?

That a system is open means not simply that it engages in interchanges with the environment, but that this interchange is an essential factor underlying the system's viability.

—Walter Buckley²

This report is intended primarily for government and nonprofit leaders who are working to create value for society and who develop projects that operate within a larger ecosystem of other units in the same organization, other organizations, and other governments and government agencies.

The concept of openness is used in a number of disciplines. In systems theory, an open system is one that continuously interacts with its environment.³ In technology, it is one that supports open standards and can therefore be made to interoperate with other computer systems.⁴ In science, it means a system that allows matter or energy to flow across system boundaries.⁵

We now talk about open source software, open data, open standards, even open education. For our purposes, the concept of “open” can be used to modify projects or organizations as well. An open project or organization tends to have porous, flexible boundaries and is receptive to contributions, resources, ideas, and direction from the outside. This is in contrast to “closed” projects and organizations which have internally defined objectives and tightly controlled resources, participation, and information flows.

Most of the organizations we work in—government and nonprofit—are traditional closed organizations, with clear-cut boundaries for employment, funding, and responsibilities. This report imagines closed organizations exploring the possibility of managing open projects that have more fluidity and exchange with their environments and with other organizations and projects. The hypothesis is that, in some situations, this approach will increase the value created for society while sharing the cost of production.

Table 1 helps explain openness by contrasting its characteristics with more commonly understood, traditional, closed concepts. Closed and open approaches are a continuum of behavior, with real-world organizations and projects being more or less open or closed on various dimensions.

2. Walter Buckley, *Sociology and Modern Systems Theory*, 1967. As quoted in W. Richard Scott, *Organizations: Rational, Natural, and Open Systems*, 2003. p. 82.

3. “Open System (Systems Theory).” Wikipedia article. http://en.wikipedia.org/wiki/Open_system_%28systems_theory%29

4. “Open System (Computing).” Wikipedia article. http://en.wikipedia.org/wiki/Open_system_%28computing%29

5. “Thermodynamic System.” Wikipedia article. http://en.wikipedia.org/wiki/Thermodynamic_system#Open_system

Table 1: A Comparison of the Characteristics of Closed vs. Open Systems

Characteristic	Closed	Open
Leadership	Clearly identified. Leadership and authority positions tightly connected.	Leadership actions can be made by a variety of actors. Leadership and authority loosely coupled.
Authority	Strong. Authority figures can firmly set and enforce direction and organize activities.	Weak. Authority figures suggest and persuade others to follow directions and engage in activities.
Membership	Clearly defined by employment, contracts, or formal declarations. Fixed, with members clearly identified.	May not be clearly defined. Dynamic, with membership determined by participation and adoption of behaviors. Members may be anonymous.
Ownership	Tightly held. Intellectual and physical property is owned by the organization or project and controlled by its authorities.	Open and shared. Intellectual property is open licensed and shared. Use is determined by users.
Boundaries	Fixed and firm. Members are either in or out.	Porous. Participants may join and leave, without giving notice.
Objectives	Unitary and clear. Established by authorities.	Different in different parts of the system and determined by participants. Changeable depending on environmental conditions.
Decisions	Made at the top of the hierarchy and passed down.	Made in many places throughout the system.
Organizational Structure	Formally planned, defined, and fixed. Rigid and hierarchical, with many relationships (e.g., employee reporting to employer); and hub-and-spoke (e.g., suppliers providing to a manufacturer)	Undefined and dynamic. Fluid networks with bidirectional and multidirectional communications and activities.
Incentives	Clearly defined, with emphasis on financial rewards—salaries, fees, bonuses.	Mixed and reliant on financial as well as other personal rewards including acknowledgement, mastery, social good.
Location	Headquarters is clearly defined. Work is done in offices and factories.	No physical headquarters. Work is done in many venues.
Work time	During formal working hours with official vacations	Happens at any time, often during vacations.

Table 2 contrasts open and closed management approaches.

The Internet, to an important extent, is the result of many separate, inter-related open projects. Managed by many different teams and organizations from around the world, these projects have developed ideas, set standards, created software, and produced products. They have improved on and built upon past projects.

Table 2: Contrasts between Closed and Open Management Approaches

Closed	Open
Tight control	Loose control
Clearly defined objectives	Fluid, emergent objectives
Enforceable deadlines	Flexible deadlines
Proprietary assets	Non-proprietary, shared assets
Secrecy expected	Secrecy discouraged
Resources are capped, tightly controlled, closely held	Resources are shared, without central control, distributed
Formal methods and processes are implementable and enforceable	Processes are harder to enforce
Superior in static settings	Superior in complex, dynamic settings
Implements known approaches	Discovers unknown approaches
Responsibility clear	Responsibility hard to identify

Open Project Approaches To Solve Big Problems

By John M. Kamensky

In this report, David Witzel describes the use of open project approaches in two arenas—the development of the Internet and the evolution of the Nationwide Health Information Network. He writes, “An open project or organization tends to have porous, flexible boundaries and is open to contributions, resources, ideas, and direction from the outside.”

Following are examples of some other IBM Center reports that reflect the principles of open project approaches in other policy arenas. These reports often refer to their approaches as cross-sector collaborative networks, but the principles are largely the same as those described in this report on open system approaches:

- ***Managing Recovery: An Insider's View***, by Edward DeSeve (2011)
The American Recovery and Reinvestment Act of 2009 provided \$787 billion of tax benefits; grants, loans, and contracts; and entitlements in more than 250 appropriation accounts across more than 25 federal agencies. The author of this report, G. Edward DeSeve, served as implementation coordinator for this law, reporting directly to Vice President Joe Biden. He created a small Recovery Implementation Office.

DeSeve used an open systems approach, which he called a managed network. The key management principles of the network included:
 - A defined network structure, with clear roles and responsibilities for participants
 - A clear common purpose among the participants
 - A governance framework agreed upon by the participants
 - Clearly delegated authority
 - Distributed accountability for action, among network participants
 - A high degree of transparency and information sharing
 - Availability of resources to get the job done
- ***Food Safety: Emerging Public-Private Approaches***, by Noel Greis and Monica Nogueira (2010)
Using food safety as a case study, the authors discuss new approaches to public-private partnerships that reflect open project principles. These public-private partnerships are seen by the authors as applicable to all complex public policy challenges, not just food safety agencies. New approaches to public-private partnerships include the following:
 - A new stakeholder model in which the private sector acts as a partner. In food safety, the private sector is acting as partner in both maintaining a safe food supply and responding to food contamination events.
 - An increased emphasis on risk-based allocation strategies. In food safety, a risk-based resource allocation will reduce disease incidence and reduce economic burden on private-sector companies that have good safety records.
 - Increased use of technology and information systems. In food safety, new food traceability techniques using private-sector information promise to speed up the recall process, thereby reducing the scale and scope of food contamination.
 - Increased use of co-regulation strategies. In food safety, co-regulation assumes a variety of forms, including setting standards, enforcement, and monitoring.
- ***Designing and Managing Cross-Sector Collaboration: A Case Study in Traffic Congestion*** by Melissa Stone, Emily Saunoi-Sandgren, John Bryson, and Barbara Crosby (2009)
In August 2007, five urban regions were selected by the U.S. Department of

Transportation to participate in the path-breaking Urban Partnership program, which was intended to integrate transit, highway pricing, technology, and telecommuting strategies aimed at reducing traffic congestion in major urban areas. The Minneapolis-St. Paul region in Minnesota was selected to receive one of the five grants. This report describes the evolution of the implementation of that initiative.

The organizers of the initiative concluded that a collaborative, multi-modal approach was crucial to making real headway on a long-standing, costly, nearly intractable public problem.

The authors describe the collaboration involved in obtaining and implementing the Minnesota Urban Partnership program. They offer insights about what contributes to successful collaboration in a collaborative network and what hinders it. The report presents lessons learned for public leaders attempting to organize collaborations, including specific lessons for project sponsors and champions.

- ***Integrating Service Delivery Across Levels of Government***, by Jeffrey Roy and John Langford (2008)

Public services are traditionally delivered through a number of government agencies via programs that are not connected to each other. In the midst of this decentralized fragmentation, two trends—a citizen-centric philosophy and a network model of service delivery—are driving demands to integrate the delivery of citizen-oriented services across levels of government. The rapid increase in technology allows this new collaborative approach to service design and delivery to be a successful substitute for the old hierarchical approach.

This report identifies strategies for, and challenges to, better integrating the delivery of citizen-oriented services. These strategies are based on experiences from Canada and four other countries that are also integrating their service delivery networks. The report concludes that at minimum, governments and their leaders need to embrace a mindset of interdependence if they hope to implement a comprehensive integration of the delivery of citizen-oriented services.

- ***Collaboration and Performance Management in Network Settings: Lessons from Three Watershed Governance Efforts*** by Mark Imperial (2004)

This report summarizes insights gained from three case studies in the management of regional watersheds and provides recommendations for all public managers operating in collaborative network settings.

The insights and recommendations presented by Professor Imperial are increasingly applicable to a wide range of public problems faced by government executives. Successful public managers are frequently finding that delivering results means working in a collaborative setting where they may have influence, but not necessarily control, over an outcome. In such an environment, managers are increasingly forming networks and partnerships to achieve objectives that no single organization or entity can achieve alone. Professor Imperial concludes, "... public managers are relying on two mutually reinforcing strategies to improve network governance: collaboration and performance management systems."

When taken together, the strategies of collaboration and performance management systems serve as powerful motivators for action and getting results in addressing complex public challenges. Professor Imperial's report provides a practical primer for all government managers on how to get started. This report will be useful and informative to all managers, including those managing our nation's watersheds.

John M. Kamensky is a senior fellow with the IBM Center for The Business of Government. He is also an associate partner with IBM Global Business Services and a fellow of the National Academy of Public Administration.

A Selective History of the Internet's Development

Today's Internet⁶ is astonishingly large and complex. By 2008, there were more devices connected to it than there were people on Earth. Estimates are that by 2020 more than 50 billion things will be connected—not just computers and phones but cars, cows, and clocks as well.⁷ Figure 1, which shows the number of people connected to the Internet, gives an idea of how fast it is growing—at a pace that approximates that of Moore's Law, doubling every two years.

It carries about 50 times more (non-spam) e-mail than the United States Postal Service does physical mail (and the USPS is the nation's second largest civilian employer, behind Wal-Mart).⁸ The majority of the Internet's infrastructure is operated by the major telecom players, and in 2009 their spending exceeded \$70 billion.

It has both created and disrupted a variety of industries, including:

- Newspapers—print circulation fell five percent in the first half of 2011 after falling more than 10 percent in the same period of 2010⁹
- Music—revenue from U.S. music sales has fallen by more than half since 1999, declining at more than eight percent each year,¹⁰ while the amount of music listened to has arguably increased.¹¹
- Video—more than 60 hours of video content is uploaded to YouTube every minute and more in one month “than the three major U.S. TV networks created in 60 years.”¹²
- Advertising—advertising on the Internet was up more than 20% in the third quarter of 2011 over the same period of 2010 and has doubled in five years.¹³
- Government—governments around the world are beginning to open up data archives and create tools for data reuse including access to salaries, donations, and lobbyist spending. Meanwhile, citizens around the world are using the power of networking to challenge more traditional power sources.

6. This paper uses a broad definition of Internet. While the term technically refers just to the network of networks that carries information around the world, in common language it has come to mean the entire ecosystem of machines, connections, data, and software that people use to send e-mail, surf the web, read news, shop, listen to music, watch movies, and more.

7. Graeme McMillan. “How Big is the Internet? Bigger than Humanity.” *Time Magazine*. July 20, 2011. <http://techland.time.com/2011/07/20/how-big-is-the-internet-bigger-than-humanity/>

8. Using the estimate of e-mail from <http://royal.pingdom.com/2011/01/12/internet-2010-in-numbers/>, USPS mail and employment from www.msnbc.msn.com/id/43259512/ns/business-us_business/t/postal-service-running-out-options/, and spam from <http://royal.pingdom.com/2011/01/19/e-mail-spam-statistics/>

9. Jeremy Peters, “Newspaper Circulation Falls Broadly But at a Slower Pace.” *New York Times*, October 25, 2010. <http://mediadecoder.blogs.nytimes.com/2010/10/25/newspaper-circulation-falls-broadly-but-at-slower-pace/>

10. David Goldman. “Music's Lost Decade: Sales Cut in Half.” *CNN Money*. http://money.cnn.com/2010/02/02/news/companies/napster_music_industry/


11. “Continued CD Sale Declines in 2008...” Music Industry News Network, March 28, 2009. http://www.mi2n.com/press.php3?press_nb=118039

12. Janko Roettgers. “YouTube's Content Explosion.” May 7, 2012. *GigaOM.com*. <http://gigaom.com/video/youtubes-content-explosion-60-hours-of-video-every-minute>

13. “Q3 '11 Internet Advertising Revenues Up 22% from Year Ago.” IAB press release, November 30, 2011. http://www.iab.net/about_the_iab/recent_press_releases/press_release_archive/press_release/pr-113011

How did we get to this point? The Internet has developed in a series of waves of innovation, each building upon the last. Here is a brief description of these waves, selectively focused on those efforts that will be used as examples in the rest of this report.

Figure 1: Growth of the Internet:1967 to 2013 (est.)



ARPANET

The commonly held notion that the original goal of the Internet was to help the world withstand a nuclear attack is a bit apocryphal, but contains a nugget of reality. One of the network's early designers, Paul Baran of RAND, was studying how communications networks could be distributed to survive destruction of centralized nodes.¹⁴

However, the original concept for starting what has become the Internet was less dramatic. It was driven by Bob Taylor's frustration at having three terminals connected to three different computers (and requiring three different sets of user commands) in his Pentagon office. Taylor was a networking visionary and the director of the Defense Department's Advanced Research Projects Agency's (DARPA) Information Processing Techniques Office. He said at the time, "There ought to be one terminal that goes anywhere you want to go where you have interactive computing. That idea is the ARPANET."¹⁵ Taylor's office launched the original ARPANET project in 1969 with the goal of connecting just four computers at four universities—UCLA, Stanford, UC Santa Barbara, and Utah (Figure 2). The computers were tied together with dedicated telephone lines provided by AT&T.

14. "Paul Baran and the Origins of the Internet." Rand Corporation, <http://www.rand.org/about/history/baran.html>

15. "An Internet Pioneer Ponders the Next Revolution." *New York Times*, December 20, 1999. <http://partners.nytimes.com/library/tech/99/12/biztech/articles/122099outlook-bobb.html> (ht: http://en.wikipedia.org/wiki/History_of_the_Internet#cite_note-2).

The contract for work on this first network was won by research and development firm Bolt, Beranek, and Newman (which later became BBN Technologies). Six months later, a two-node network connecting UCLA to Stanford was tested with the message “login.” The system crashed after the letters “l” and “o” were transmitted, requiring the first of what have been many generations of Internet fixes.¹⁷


By the end of 1969, ARPANET was able to carry commands and data. A small crowd of programmers from universities and a few research institutes around the country went to work developing new software and testing ideas. Protocols to support basic functions were quickly developed. File transfer (FTP) and data entry (Telnet) protocols were released in 1971 and the now-ubiquitous @ was selected to define e-mail addresses.

Also in 1972, BBN and DARPA organized the first International Conference on Computer Communication at D.C.’s Washington Hilton to show off the nascent network and its novel applications. Under the guiding hand of BBN’s Bob Kahn, director of the original ARPANET design project, this first-of-its-kind event went off with only minor glitches, enough to convince academic, government, and industry representatives that ARPANET was viable. By 1975, with about 100 computers connected, the network was established enough that DARPA could hand off daily operations to the Defense Communications Agency.

In addition to wires being connected to computers, the Internet’s do-it-ourselves culture was being established as well. Many of the proposals and decisions that define today’s Internet were raised in Requests for Comment (RFCs).¹⁸

As RFC 1549 explains, RFCs “are working notes of the Internet research and development community. A document in this series may be on essentially any topic related to computer communication, and may be anything from a meeting report to the specification of a standard.”¹⁹

Figure 2: Diagram of the Original, 4-node, ARPANET¹⁶



Understanding Requests for Comments (RFCs)

Some RFCs have more practical application than others. Part of a tradition of April Fool’s Day contributions, RFC 1149 describes how to transmit Internet traffic via carrier pigeon. It was recently updated, in RFC 6214, to support new generation, IPv6 addresses. Strongly suggesting that people do indeed have time to help solve complex problems, the carrier pigeon protocol has been successfully tested in the field.²⁰

16. “Early sketch of ARPANET’s first four nodes.” *Scientific American*, December 4, 2009. http://www.scientificamerican.com/gallery_directory.cfm?photo_id=5B11E498-A639-3996-6D74347AFB957CAA

17. This section relies heavily on a set of resources including Wikipedia’s article on ARPANET (<http://en.wikipedia.org/wiki/ARPANET>); Katie Hafner and Matthew Lyon’s *Where Wizards Stay Up Late* (1996), a story of the development of the Internet up to the early 1990s; and Vinton G. Cerf’s “Computer Networking: Global Infrastructure for the 21st Century.” (<http://www.cs.washington.edu/homes/lazowska/cra/networks.html>)

18. “Request for comment” is really a misnomer as the documents rarely request comments and more often explain or define.

19. RFC 1594: “FYI on Questions and Answers: Answers to Commonly Asked ‘New Internet User’ Questions,” <http://tools.ietf.org/html/rfc1594>

20. “IP Over Avian Carriers.” Wikipedia article. http://en.wikipedia.org/wiki/IP_over_Avian_Carriers

A Selective List of Pioneers of the Internet

Marc Andreessen	While working at the National Center for Supercomputing Applications (NCSA) at the University of Illinois, Andreessen was a primary developer of the Mosaic web browser, released in 1993. Mosaic helped popularize the World Wide Web by displaying pictures in line with text. Andreessen, now a venture capitalist, was in the first round of inductees to the World Wide Web of Hall of Fame.
Brian Behlendorf	Computer programmer and open source software guru. He led development of the Apache™ Web server and served as President of the Apache Foundation. He co-founded software tools company CollabNet, which sponsors Subversion®, an open source software version control system. He is now Chief Technology Officer for the World Economic Forum.
Yochai Benkler	Harvard professor, author, and director of the Berkman Center for Internet and Society. His writings, including his influential paper “Coase’s Penguin,” helped define and explain the open source software development and what he calls peer production.
Tim Berners-Lee	English-born computer scientist and professor is best known as the inventor of the World Wide Web, for which he was knighted in 2004. He is now director of the World Wide Web Consortium, which oversees development web standards.
Vint Cerf	American computer scientist, often called the father of the Internet for his role in developing the TCP/IP protocol along with co-father Bob Kahn, while they were at DARPA. He helped found ICANN, which oversees internet addressing and helps maintain its operations. One of his current efforts is to launch the interplanetary Internet.
Steve Crocker	Computer scientist and professor. Crocker was one of the first Internet developers while at student at UCLA and author of the first Request for Comments (RFC). He helped create the original Network Working Group and set the tone of a do-it-ourselves Internet. He was in the first class inducted into the Internet Hall of Fame in 2012.
Bob Kahn	Engineer and computer scientist who shares, with Vint Cerf, title of “father of the Internet.” Kahn was on the original team at BBN to develop the first ARPANET equipment and software and worked with Cerf to define TCP/IP. In between he helped organized the first International Conference on Computer Communications (ICCC).
Robert Metcalfe	Electrical engineer best known for helping inventing Ethernet. A graduate of MIT and Harvard, he was involved in connecting MIT computers to ARPANET in the early 1970s. He coined Metcalfe’s Law: the value of a communications network is proportional to the square of the number of connected users.
Jarkko Oikarinen	Finnish computer programmer developed Internet Relay Chat (IRC), the first Internet chat network. He currently works in Sweden for Google.
Jon Postel	Computer scientist and editor of the Request for Comment series for 29 years until his death in 1998. He was involved in early work on ARPANET while he was a student at UCLA. He is known for Postel’s Law which says “be conservative in what you do, be liberal in what you accept from others.”
Eric Raymond	Computer programmer and author. Unofficial spokesperson for the open source movement and author of <i>Cathedral and the Bazaar</i> .
David Reed	Computer scientist and Adjunct Professor at the MIT Media Lab. He was involved in early development of TCP/IP and is known for Reed’s law, which states that the utility of large networks can scale exponentially with the size of the network.
Brian Reid	Computer scientist from University of Maryland and Carnegie Mellon. While at Carnegie Mellon he was part of the team defining Internet e-mail standards.
Richard Stallman	Computer programmer and activist. Launched the GNU (for “GNU not Unix”) Project in 1983 and founded the Free Software Foundation. He is the main author of the General Public License (GPL), a widely used free software license.
Robert (Bob) Taylor	Visionary who led initial development of ARPANET while Director of the Information Processing Techniques Office (IPTO) of the Advanced Research Project Agency (ARPA). He later ran the technology research center at Xerox PARC.
Ray Tomlinson	Computer programmer famous for sending the first network e-mail and choosing the @ sign to demarcate e-mail addresses. He worked at BBN and helped develop the original ARPANET protocol NCP as well as TELNET. He was inducted into the Internet Hall of Fame in 2012.
Linus Torvalds	Finnish computer programmer famous for starting development of the open source Linux operating system. He is currently Fellow at the Linux Foundation and still the ultimate authority on what new code is incorporated into the standard Linux kernel. He was inducted into the Internet Hall of Fame in 2012.
Jimmy (Jimbo) Wales	Entrepreneur and cofounder of Wikipedia with the vision of “a world in which every single person on the planet is given access to the sum of all human knowledge.” He is on the Board of Trustees of the Wikimedia Foundation.
Dave Winer	Software developer, author, and entrepreneur. An advocate for web syndication, he helped develop Real Simple Syndication (RSS) and turned over ownership of the format to Harvard University. He also coined the term “narrate your work,” advice he would give staff at his firm UserLand Software.

The first RFC,²¹ addressed to the “Network Working Group” in 1969, was prepared by Steve Crocker while he was at UCLA. It dealt with the specifics of how the new computer called an interface message processor or IMP (the first router), being developed by BBN, would communicate with university host computers which were to become the first nodes on the soon-to-launch ARPANET.

More than 6,500 RFCs have been issued since 1969. Though early RFCs were distributed via U.S. mail, distribution now is primarily via e-mail, newsgroups, and more recently, the web, relying on the infrastructure the RFCs defined. Organization of RFCs is managed by the RFC editor, a role first held by Jon Postel for 29 years until his death in 1998.

The addressee of the original RFCs, the Network Working Group, isn’t a formal group in the sense we think of for businesses or government. According to *RFC Editor*, “Network Working Group” meant the set of researchers who developed the packet switching protocols for the ARPANET and now “should be interpreted as the set of users, vendors, and researchers who are working to improve and extend the Internet.”²² Now RFCs are addressed to the Internet Engineering Task Force (IETF), the equally informal successor to the Network Working Group. The IETF doesn’t have an official charter and operates as a loose organization of kindred spirits.²³

The Network of Networks

While ARPANET had moved from a research project to operations, it was only one network of what were beginning to be many. It represented one stage in a broader connectivity vision that included linking many stationary networks, satellites, and naval ships; today, trains, planes, and even spaceships have been connected.²⁴

Vint Cerf, then at Stanford, and Bob Kahn, organizer of the successful International Conference on Computer Communication, who had since moved to DARPA, tackled the problem of linking networks together. In 1974 they released a paper describing a network protocol called TCP.²⁵ By 1977 this protocol was in place and used to connect ARPANET with a satellite network and a radio network, creating a true inter-net. The protocol was further refined, including being broken into two pieces—TCP and IP—which were described in RFCs 760 and 761, released in 1980. TCP/IP protocols define the modern Internet.

New services continued to be added to the network, with a particular focus on those that supported conversation and content sharing. E-mail, one of the oldest and still most widely used communications technologies, predates ARPANET, though at first it could only be used between terminals on the same host computer. Ray Tomlinson, another BBN alumnus, is credited with making changes to an existing e-mail program and sending the first networked e-mail in 1971, one that went from a user on one host computer to a user on another computer. To identify mailboxes on different computers, he chose the now-famous @ sign to sit between the user’s name and the user’s computer’s name. He explained it “just made sense.”²⁶

21. Steve Crocker. RFC #1, April 7, 1969, <http://tools.ietf.org/html/rfc1>

22. *RFC Editor* FAQs. <http://www.rfc-editor.org/rfcfaq.html#wg>

23. Jeff Doyle. *OSPF and ISIS: Choosing an IGP for Large-Scale Networks*. 2005. <http://fengnet.com/book/OSPFandISIS/ch01lev1sec6.html>

24. “Interplanetary Internet.” Wikipedia article. http://en.wikipedia.org/wiki/Interplanetary_Internet

25. Vinton Cerf and Robert Cahn. “A Protocol for Packet Network Intercommunication.” Paper. <http://ece.ut.ac.ir/classpages/F84/PrincipleofNetworkDesign/Papers/CK74.pdf>

26. Ray Tomlinson. “The First Network E-mail.” <http://openmap.bbn.com/~tomlinso/~/ray/firste-mailframe.html>

Since then e-mail has been through many generations of improvement. Dave Crocker recalls²⁷ that by 1975, many others had built on Ray Tomlinson's work, including Lawrence Roberts, Barry Wessler, Marty Yonke, John Vittal, and Crocker himself. Over time, enhancements have added text editors to help write messages, folders to sort them, the ability to forward e-mail and reply to all recipients, ways to download and read e-mail offline, and, eventually, the ability to integrate and attach pictures, documents, and other files. The emergence of standards that allow users to attach and e-mail files to each other, and that allow networked computers to transfer e-mail reliably, greatly extended the value and reach of the tool.

An early decision to break the e-mail problem into two pieces—placing details describing where to send a message into a “header” and content of the message in the “body”—allowed for rapid improvement. This modularization allowed the two components to evolve independent of one another. For example, SMTP was developed to improve e-mail transport and MIME was developed to support e-mailing various types of content (e.g., pictures, Word files, spreadsheets).

But there were other communications channels developing as well. For example, newsgroups facilitating discussion on topics were launched in 1979 by Duke University students. Echoing the evolutionary development of many other Internet tools, the original system developed at Duke (called *A News*), was improved at Berkeley with *B News* and then again at University of Toronto with *C News*.²⁸

The most famous implementation of newsgroups was through *Usenet*. Articles could be posted to different topics and then copied to host computers around the world via a variety of mostly dial-up approaches. By 1985 a new protocol, NNTP, was developed to take better advantage of the spread of the Internet. Today there are over 9,000 active newsgroups,²⁹ though the importance of newsgroups and NNTP has declined as they have been supplanted by the World Wide Web.

Another communications tool, Internet Relay Chat or IRC, allowed groups of network users to type back and forth to each other. Created in 1988 by Jarkko Oikarinen, a programmer from Finland, IRC is used to host real-time chat rooms and global discussions.

Gopher, an early method of sharing documents and allowing them to link to one another, was released in 1991. It was named by its University of Minnesota development team as a play on the university's mascot, the Golden Gopher. Development of WAIS, or wide area information servers, began in the early 1980s and continued for much of the decade; an open source version was funded by the National Science Foundation. WAIS, a text searching system, often supplemented Gopher servers to help with information retrieval.³⁰

By 1984 about 1,000 hosts made up a true Internet that connected computers and networks at U.S. universities, government agencies, and defense facilities with those in Europe and Asia. ARPANET was fully decommissioned by 1989, replaced by a TCP/IP network with backbone connections funded by the National Science Foundation.³¹

27. Dave Crocker. “E-mail History.” <http://www.livinginternet.com/e/ei.htm>

28. “Usenet.” Wikipedia article. en.wikipedia.org/wiki/Usenet

29. Harley Hahn. “Usenet Tutorial.” *Harley Hahn's Usenet Center*. <http://www.harley.com/usenet/usenet-tutorial/newsgroups-and-hierarchies.html>

30. “Wide Area Information Server.” Wikipedia article. http://en.wikipedia.org/wiki/Wide_area_information_server


31. Walt Howe. “A Brief History of the Internet.” *Walt's World Wide Webbing*. <http://www.walthowe.com/navnet/history.html>

The World Wide Web

The cable had been laid and computers connected, but the Internet was still largely a tool for computer software and hardware engineers and, to a lesser extent, adventurous researchers at universities and laboratories. In 1992 commercial subscribers, in addition to government and academic staff, were allowed access and services like CompuServe began to offer gateways for their customers to connect to the Internet. Nonetheless, it was an information superhighway with few vehicles.

That changed with adoption of the World Wide Web. Tim Berners-Lee and collaborators at the international research institute, CERN, created the definitions and developed the first generation of tools needed to build and browse a website. HyperText Transfer Protocol (HTTP) and HyperText Markup Language (HTML) defined how the web would work, while the first web browser, editor for making web pages with HTML, and server software to respond to requests for web pages from the web browser demonstrated how it would work.³² Berners-Lee announced the project in public (via a newsgroup) in August 1991, offering the code and details for free and asking for more collaborators.³³

The working model and willingness to share motivated others to contribute and extend upon the ideas from CERN. The National Center for Supercomputing Applications (NCSA), a project at the University of Illinois, released a new web browser, Mosaic, in 1993. Mosaic ran on PCs and added the ability to display pictures. Its lead author, Marc Andreessen, went on to help found Netscape Communications and grace the cover of Time Magazine.³⁴



NCSA also released a popular web server which became the basis for further development by a group calling itself the Apache Group, coordinated with the help of Brian Behlendorf. Apache quickly became and remains the most popular web server on the Internet. Behlendorf was a primary developer of the software and later became president of the Apache Software Foundation, which manages the web server's code and development.

Thanks to these enhanced tools and the creators of websites around the world, Berners-Lee's creation surpassed Gopher, WAIS, and other content-sharing strategies to dominate the Internet. Now, 30 years later, there are more than 200 million websites.³⁵

Another important project that was supported by and now supports the Internet was also born in 1991, this time in Finland. A young programmer named Linus Torvalds set out to get a better understanding of the workings of the new Intel 386 computer processor by developing an operating system to run on it. In his initial e-mail to perspective contributors, Torvalds explained "I'm doing a (free) operating system (just a hobby, won't be big and professional ...)."³⁶

His hobby has become the open source, collaboratively written operating system called Linux, a cornerstone of the Internet itself. It has grown to create its own market segment worth tens

32. "World Wide Web." Wikipedia article. <http://en.wikipedia.org/wiki/Www>

33. Tim Berners-Lee. "Message from discussion 'Qualifiers on Hypertext Links.'" Google Groups discussion. August 6, 1991. <http://groups.google.com/group/alt.hypertext/msg/06dad279804cb3ba>

34. "Netscape's Marc Andreessen." *Time* magazine cover. February 19, 1996. <http://www.time.com/time/covers/0,16641,19960219,00.html>

35. Amit Agarwal. "How Many Websites Are There on the Internet?" *Digital Inspiration*. November 4, 2008. <http://www.lanlol.org/internet/total-websites-on-internet-worldwide/5206/>

36. "The First Linux Announcement from Linus Torvalds." *The Linux Daily*. <http://www.thelinuxdaily.com/2010/04/the-first-linux-announcement-from-linus-torvalds/> (ht: Clay Shirky, *Here Comes Everybody*, p. 238.)

of billions of dollars. Linux is made up of over 14 million lines of computer code (roughly equivalent to 1,100 200-page books) and is being updated at the rate of five changes per hour.³⁷ As Ryan Paul explains, “It runs a majority of the world’s supercomputers, many of the most popular websites, and a large number of mainstream consumer electronics products. If you search the Web with Google, interact with your friends on Facebook, look up a fact on Wikipedia, read a book on a Kindle, watch Netflix on a Boxee Box, or make a call with an Android phone, you are relying on Linux.”³⁸

Wikipedia was one of the first high-profile projects to take the Internet’s do-it-ourselves mantra beyond computing. It is built upon an open source software tool called a wiki which was developed by Ward Cunningham in the mid-1990s and named from the Hawaiian word for quick. Wikis allow multiple users to add and change content on a website while tracking the history of these changes and allowing easy return to past versions.

Wikipedia used this easy-to-change approach to build the world’s largest and most-used encyclopedia, with over 20 million articles and an average of 25 million visitors a day to its English-language site. Launched in 2001, its success builds from a very specific self-image, that of an encyclopedia with a neutral point of view.

Web 2.0

The term Web 2.0 was popularized by pundit and technology storyteller Tim O’Reilly. It summarized a set of changes happening on the web that encouraged information sharing between websites and relied heavily on contributions from web users. The focus of the web was evolving from an old-media model of products being delivered to consumers (e.g., Hollywood producing a movie for customers to watch) to a new approach of products being created by many participants and shared (e.g., users uploading movies to YouTube for each other to watch, comment on, and mimic). In the Web 2.0 world easy sharing leads to very rapid creation, reuse, and revision as participants build upon and riff off of each other’s ideas.

While a fundamental component of Berners-Lee’s web concept was the ability to organize and reuse materials via links, options for reuse became more sophisticated as the ability to share content and services was built into web software. RSS, or real simple syndication as it was dubbed by one of its designers, Dave Winer, became a common feature in blogs and websites. RSS provided a standard format so one website could reuse content from another or aggregate content from multiple websites. For example, the *New York Times*’ RSS feed meant other websites could automatically feature *New York Times* headlines.

Application programming interfaces (APIs) also became increasingly common. APIs allowed web developers to build from services provided by other websites. A very popular API is the one Google provides for Google Maps that allows maps to be personalized and embedded, resulting in the maps seen on websites advertising restaurants and vacation houses around the world.

At the same time that features were being added on the software front, lawyers were innovating on the content side. Computer activist and programmer Richard Stallman popularized the idea of open licenses for software. In 1989, his GNU³⁹ General Public License (or GPL) was

37. Ryan Paul. “March of the Penguin: ARS Looks Back at 20 Years of Linux.” *Ars Technica*. August 25, 2011. <http://arstechnica.com/open-source/news/2011/08/march-of-the-penguin-ars-looks-back-at-20-years-of-linux.ars/3>

38. Ibid.

39. GNU is a recursive, geeky joke meaning “GNU, not Unix” to distinguish Stallman’s products from the Unix software, whose license was owned by AT&T.

released as a general license for any software product. It allowed programmers to reuse and repurpose software code as long as the resulting product still carried the GPL license. This popular open license is used by many common pieces of Internet infrastructure, including the Linux kernel and the MySQL database. In 2010, the White House's web team released software they developed under this license.⁴⁰

In 2002, Creative Commons extended the GPL concept and released a set of copyright licenses that provided legal language to allow content creators to license their works for free to the public. These licenses provided a set of options for what uses the content could be put to and what citations were required. For example, a photograph may be licensed so that it can be reused in any way as long as the owner is credited or so that it can only be used for non-commercial purposes. Through services like its handy license picker and getting the licenses ported to over 50 jurisdictions around the world, Creative Commons has generated wide adoption, claiming over 350 million licensed works.⁴¹ In 2009, the Wikimedia Foundation adopted a Creative Commons license for Wikipedia.

The Next Wave

It is likely that these waves of innovation will continue, with the one now underway characterized by a move to cloud computing and mobile devices. Cloud computing, affordable computing services provided by commercial organizations like Amazon and Rackspace that are being gradually implemented in the federal government, gives innovators and entrepreneurs access to massive amounts of storage and computing power while paying for only what they use. This allows them to affordably start small projects that can grow very large and handle surges in attention.

The move to mobile devices, including smartphones and tablet computers, has increased the number of Internet users to include almost all phone users, and has dramatically expanded the Internet's footprint into new geographies, businesses, and roles.

40. Rob Winokates. "WhiteHouse.gov Releases Second Set of Open Source Code." February 11, 2011. White House blog. <http://www.whitehouse.gov/blog/2011/02/11/whitehousegov-releases-second-set-open-source-code>

41. History. Creative Commons webpage. <http://creativecommons.org/about/history>

Lessons from the Internet

The Internet has been a fertile test bed for experimentation and innovation leading to new products, new approaches to communication and collaboration, and new organizational and ownership structures. To a large extent, what we know as the Internet today is the result of many individual open projects, often run by groups of volunteers. These projects built new services or improved upon existing services and used existing interfaces in new ways. The process unleashed creativity and innovation, reached large scale, and had a big impact on society.

Government and nonprofit leaders often want to achieve these same outcomes in their projects. However, to be successful they need several things to happen. They need more coordination among overlapping efforts to reduce redundancy and increase progress. They need to accelerate learning to avoid reinventing the wheel and to improve quality and performance. They need to increase participation and encourage ownership to expand scale and accelerate adoption. They need to foster creativity and innovation to solve sticky problems and they need approaches that evolve and continually improve.

The dozen tips that follow extract lessons from the creation and evolution of the Internet to help government and nonprofit leaders spur more innovation, increase their reach, and have a larger impact.

Tip One: Let Everyone Play

Relax, all right? Don't try to strike everybody out. Strikeouts are boring! Besides that, they're fascist. Throw some ground balls—it's more democratic.

—Crash Davis to Nuke LaLoosh in *Bull Durham*, 1988

Collaborators welcome!

—Tim Berners-Lee, August 1991⁴²

The Internet has grown up with a strong do-it-ourselves culture that is open to contributions from many quarters and reliant on its own participants for ideas, fixes, and improvements. The IETF, the organizing entity of Internet standards, is an informal working group open to anyone who wants to participate. Important pieces of today's Internet, including the World Wide Web, Linux, Apache, and Wikipedia, were started by small teams that expanded into broader, let-everyone-play collaborations. Opening up projects to more participants helps bring in additional resources, attracts advocates, and captures new ideas.

42. Tim Berners-Lee. "Message from discussion 'Qualifiers on Hypertext Links.'" Google Groups discussion. August 6, 1991. <http://groups.google.com/group/alt.hypertext/msg/06dad279804cb3ba>

A variety of arguments are made to explain the value of creating an inclusive problem-solving process. Bill Joy, one of the founders of Sun Microsystems, allegedly said “No matter who you are, most of the smartest people work for someone else.”⁴³ Similarly, Eric Raymond argues that including more people makes it easier to solve problems because, “Given enough eyeballs, all bugs are shallow.”⁴⁴

Perhaps most important, openness to contributions from many sources provides new opportunities for innovation and creativity—what Jonathan Zittrain calls generativity. He defines generativity as “a system’s capacity to produce unanticipated change through unfiltered contributions from broad and varied audiences.”⁴⁵ This ability to discover new approaches gives open projects a special edge over traditional, closed teams.

Engaging a wide population of potential contributors can improve innovation and problem solving. It can help a project reach beyond what it knows is possible, beyond its perception horizon, to learn from people and experiences on the edge (Tip Six). The Internet’s do-it-ourselves inclusiveness has meant it has attracted contributions from many countries of the world, including the examples of the U.S., U.K. Switzerland, and Finland in the preceding brief history. Likewise, it has benefited from contributions from many organizations—Department of Defense, CERN, BBN, UCLA, and Google, to name just a few.

Letting everyone play requires special leadership and management approaches. The imagination, management, and leadership required to make dynamic distributed teams successful is a vital resource. Brian Behlendorf, one of the creators and custodians of the Apache web server, writes, “The brilliance inside Linus [Torvalds] is his ability to orchestrate the aggregated interests of thousands of other developers, all individually scratching their own itch (or that of their employer), and thereby making a product renowned for reliability, performance, and the features people need.”

Replicating Torvalds’ brilliance is not an easy thing. Most of the following tips are intended to help create opportunities for wider participation and use those new resources productively.

Tip Two: Play Nice

When you read RFC 1, you walked away from it with a sense of, “Oh, this is a club that I can play in too.”

—Brian Reid⁴⁶

Professor and computer scientist Brian Reid recalls the first RFCs and the Network Working Group, saying, “It has rules, but it welcomes other members as long as the members are aware of those rules ...The language of the RFC was warm and welcoming. The idea was to promote cooperation, not ego.” In RFC 1000,⁴⁷ Crocker explains the setting for his early decisions. “Most of us were graduate students and we expected that a professional crew would show up eventually to take over the problems.” This lack of official authority motivated the inviting tone. “I remember having great fear that we would offend whomever the official protocol designers were, and I spent a sleepless night composing humble words for our notes.”

43. “Joy’s Law.” Wikipedia article. http://en.wikipedia.org/wiki/Joy%27s_Law_%28management%29

44. Eric S. Raymond. *The Cathedral & the Bazaar*, 1999, p. 19.

45. Jonathan Zittrain, *The Future of the Internet and How To Stop It*. 2008, p. 70.

46. Katie Hafner and Matthew Lyon. *Where Wizards Stay Up Late*. 1996, p. 144.

47. J. Reynolds and J. Postel. RFC 1000. August 1987. <http://www.ietf.org/rfc/rfc1000.txt>

Of course, there was no professional crew in the wings and Crocker and the informal Network Working Group continued to define the Internet, creating the rules of the road as they went along, exhibiting the do-it-ourselves mentality that characterizes Internet innovation. The RFC's "this is a club that I can play in too" tone lives on and creates space for letting everyone play (Tip One). Ease of entrance is an important component in collaborative projects, especially when you have no personal or institutional ties to the people you are collaborating with.

To make it easy for other people to play, rules and processes that would keep people out or scare them away should be avoided. A common mistake in collaborative projects is to inadvertently reduce participation by setting high walls to protect the project from low-quality contributions or vandalism. Often these protections make it more difficult to participate, and the potential problems were either unlikely to happen or inexpensive to fix if they did happen. As an alternative, Jonathan Zittrain suggests his procrastination principle: "waiting for problems to arise before solving them."⁴⁸

Jimmy Wales, founder and benevolent dictator of Wikipedia, explains the issue this way. "If you prevent people from doing bad things, you prevent them from doing good things." Wales illustrates the implications of solving problems prematurely with a restaurant metaphor. "If you serve steak, that means steak knives, which are really dangerous in the wrong hands, so you need to put barriers between the tables ..."⁴⁹ How many people will visit a steak house with armored cubicles around each table?

Instead of trying to guard against the damage users might do, Wikipedia makes it easy to recover from damage that has been done. A standard component of the wiki is the ability to revert to a past generation of a page. Thus, should a page be damaged, intentionally or not, it can quickly be returned to its previous state. Increasingly complex tools help Wikipedia contributors track and respond to problems.

Playing nice means more than making it easy to participate. Good open project design will deliberately take advantage of participants' personal incentives. The success of many open source software efforts has convincingly demonstrated that programmers (at least) will contribute to joint projects even if they aren't paid by the project. As we've seen in the case of Linux, they may still be paid, especially when their employer has an interest in development, but often projects start as labors of love. It isn't just software developers. We are seeing unpaid contributors editing articles for Wikipedia, writing book reviews for Amazon, and identifying craters on Mars for NASA.⁵⁰

Part of what drives this willingness to contribute is a change in the cost structure of contributing. Because of affordable, powerful, widely available computers (and now smartphones) with Internet connections, people are able to make contributions that potentially affect millions at low cost to themselves. Internet visionary Clay Shirky calls this the "half-life of love," explaining, "now we can do things for strangers who do things for us, at a low enough cost to make that kind of behavior attractive, and those effects can last well beyond our original contribution."⁵¹

The other realization is that people have always wanted to contribute and now they have more ways to. Daniel Pink summarizes the state of understanding about motivations in his book *Drive*, saying, "When it comes to motivation, there's a gap between what science knows and what business does. Our current business operating system—which is built around external,

48. Zittrain, p. 135.

49. <http://longnow.org/seminars/02006/apr/14/vision-wikipedia-and-the-future-of-free-culture/> (requires registration and fee)

50. <http://en.wikipedia.org/wiki/Clickworkers>

51. Shirky, p. 142.

carrot-and-stick motivators—doesn't work and often does harm ... This new approach has three essential elements: (1) Autonomy—the desire to direct our own lives; (2) Mastery—the urge to get better and better at something that matters; and (3) Purpose—the yearning to do what we do in the service of something larger than ourselves.”⁵²

Leadership in open projects implies creating space for participants to join in decision making, build skills, make contributions, and find meaning. Project designers who successfully engender these characteristics can hope to attract motivated participants.

Tip Three: Talk About What You Are Doing While You Are Doing It

The public servant needs to presume to publish.

—Nigel Shadbolt⁵³

Open projects don't have the internal coordination tools a traditional team in a business or government agency has. In fact, since many participants and potential participants may not have even identified themselves, it can be impossible to assemble or coordinate the whole team. Thus, participatory projects need to find open approaches to coordinate activities.

Innovator and software developer David Winer coined the term “narrating your work” and suggests it as one solution. He used the term in his distributed software company to mean talking about what you are doing while doing it—what problems are you having, what are your priorities, who are you working with, what alternatives did you consider and discard, what worked, and, perhaps more important, what didn't work.⁵⁴ This is in contrast to traditional communications approaches that focus on internal discussion of operations and public announcement of only (usually positive) results.

Public sharing of project details creates three types of value. First, if you are letting other people play it will improve coordination with the projects and people who might want to play. It allows outside participants and potential participants to listen in to understand objectives, suggest improvements, pick opportunities to make contributions, and decide whether they want to play more actively.

Second, it also allows outside people and projects to adjust their activities to take advantage of your learning, avoid overlap with your activities, and create adjacent value. Clay Shirky explains how this can work:

This is the secret of the open source ecosystem and, by extension, of all the large-scale and long-lived forms of sharing, collaborative work, and collective action now being tried. Because anyone can try anything, the projects that fail, fail quickly, but the people working on those projects can migrate just as quickly to the things that are visibly working. Unlike the business landscape, where companies have an incentive to hide both successes (for reasons of competitive advantage) and failures (to forestall any perception of weakness), open source projects advertise their successes and get failure for free.⁵⁵

52. Daniel H. Pink, *Drive: The Surprising Truth About What Motivates Us*. 2009, pp. 203–204.

53. Rosalie Marshall. “Open Data Tsar Calls for Change in Government Mentality.” *V3.co.uk*. February 8, 2012. <http://www.v3.co.uk/v3-uk/news/2144927/tsar-calls-change-government-mentality>

54. David Winer. “Narrate Your Work.” *Scripting.com*. August 9, 2009. <http://scripting.com/stories/2009/08/09/narrateYourWork.html>

55. Clay Shirky, *Here Comes Everybody*, 2008, p. 258.


It is important to note that much of this knowledge-sharing is directly from participant to participant without the need for a central moderator or editor. Author John Hagel explains that “knowledge flows on the edge” (see Tip Six) are particularly valuable. It is this information that helps expand the perception horizon and engage participants who “focus on ways to innovate and create value by connecting unmet needs with unexploited capabilities and then scaling these opportunities as rapidly as possible.”⁵⁶

The third benefit to narrating work is trust-building. Transparent discussion of what is being done and why makes it easier for potentially competing groups to understand, trust, and participate. The architect of the HHS NwHIN Direct Project (discussed further on pages 42–44), Arien Malec, explains that private documentation can create divisions between those in the know and those outside. These divisions can create dissent, engender opposition, and reduce participation.

Tip Four: Use Multiple Channels of Communication

*It's no use, Mr. James, it's turtles all the way down.*⁵⁷

While the original intent of ARPANET was to share expensive hardware resources by linking computers together, the side effect of linking *people* together has probably turned out to be more important. An early focus for Internet developers was to create and improve on a variety of communications tools. The tools had many different characteristics, including providing users options of how many people could participate at a time, whether participants had to be online at the same time, how easy it was to keep track of discussions, how hard it was to archive content, whether participants could be anonymous, whether access could be restricted to specific group members, how hard it was for participants to find each other, and how easy it was to reuse materials.



Growing selection and sophistication of communication tools was generated by many rounds of improvement to the original tools including e-mail, newsgroups, and chat—think of the incremental progression from A News to B News to C News (also see Tips One and Seven)—plus the addition of many new modalities like Inter-Relay Chat (IRC), instant messaging, blogs, Skype, Facebook, Twitter, and more recently Google+.

By using multiple channels of communication, Internet pioneers were able to support the complexity of having lots of people involved in a project. They tended to use overlapping channels of communication that fill different niches and reach different readers, allowing them to find interested participants and keep them current and coordinated.

According to Behlendorf, the Apache Project has an unwritten policy to use e-mail lists for discussions about architecture and when consensus is needed. He explains that e-mail lists are

56. Hagel. p. 53

57. John S. Wilkins. “Turtles All the Way Down.” *Evolving Thoughts*. March 28, 2011. <http://evolvingthoughts.net/2011/03/turtles-all-the-way-down/> (graphic from <http://somethinghappens.keenspot.com/d/20070531.html>)

better at providing “traceability and sense of history that communities and decisions need.” However, they use real-time IRC for specific purposes, including during global member meetings. He says it is “great for answering newbie questions, developers coordinating on a release or bughunting, building a sense of team spirit, etc.,” but can be noisy and hard to search.⁵⁸

Much critical discussion and documentation are embedded in the actual development tools that support projects. The open source community has led development of its own software code repository, Apache™ Subversion® (another Behlendorf contribution). The repository stores the current generation of a software product as well as past generations of code (allowing software to be rolled back to past versions), documentation of changes, and notes about what changes have been made, why, and by whom.

Wikipedia relies heavily on the revision history built into wikis. Much of its content lies behind the View History and Talk tabs on every page. Every change to every article is tracked in its history while the Talk tab allows editors and browsers to announce the current status of a page, discuss the rationale and approach to a particular topic, and question, explain, or plan changes.

The multitude of discussion channels and volume of discussion supports distributed problem solving in a number of ways.

- It provides documentation and historical perspective on the progress that has been made and reasons for decisions.
- Interested but inexperienced newbies can learn about the project.
- Project team learning accelerates as participants find each other and share their tacit knowledge, sharing what John Hagel calls know-how rather than know-what.⁵⁹
- Related projects can find each other, making it easier to coordinate or rationalize efforts.
- Feedback helps to identify problems and focus and refocus efforts.
- Other projects benefit when “leaky” systems create information spillover, wherein, as author Steven Johnson says, “good ideas have a tendency to flow from mind to mind.”⁶⁰

Tip Five: Give it Away

After the music leaves us, it's on its own. If they want to collect it and show it around or re-live it, or whatever, that's okay with me.

—Jerry Garcia⁶¹

One of the truly amazing characteristics of the Internet is that much of the key intellectual property that makes it run is available in the public domain for anyone to use and reuse. Concepts like TCP/IP and HTML, software like the Linux Operating System and Apache Web Server, and content like that of Wikipedia are available under open licenses like GPL and Creative Commons. Using public domain intellectual property as Internet infrastructure isn't a rule (in fact, Internet architecture doesn't have enforced rules) though it is a recommended practice.

Depending on widely available and reusable intellectual property has a number of positive benefits. It is a key component of getting others to play (see Tip One). The opportunity for

58. Author communications with Brian Behlendorf.

59. John Hagel III, John Seely Brown, Lang Davison. *The Power of Pull*, 2010. p. 54.

60. Johnson, p. 53.

61. “Grateful Dead: Jerry Garcia Speaks, Volume 1.” *NBC News on Stage*. <http://www.hulu.com/watch/11363/nbc-news-on-stage-grateful-dead-jerry-garcia-speaks-vol-1> (ht: Chris Bucchere)

people to learn from public knowledge assets and be sure of their right to leverage their contributions later is a great incentive for participation.

It also tends to increase the social return from a project. Way back in 1994, *Wired* magazine founder Kevin Kelly suggested we “cultivate increasing returns.”⁶² Kelly said, “each time you use an idea, a language, or a skill you strengthen it, reinforce it, and make it more likely to be used again.” Ideas, algorithms, standards, and software are all inexpensively reproduced for distribution and improvement, generating increasing returns to our efforts.

Open licensing can help projects create a “self-leveling playing field.”⁶³ Government and non-profits are often wary of being responsible for picking winners and losers. Open knowledge sharing creates a shared asset that all competitors can benefit from in a pre-competitive space. The flip side of not wanting to be responsible for picking winners is that governments and nonprofits don’t have to be proprietary about what they know and learn and are thus able to “lift all boats” with shared knowledge and resources.

Understanding a Fork

A stabilizing characteristic of open source projects is the threat of a fork, which happens when one or more developers split the development team, take a copy of the source code, and start a different path. As Eric Raymond explains, a fork is generally considered a bad thing.⁶⁴ With most open source licenses, this is legal and can be done without permission. A fork is a form of voting with your feet, and its threat helps encourage comity and compromise among the development team.

In 2009 a prominent open source database package, MySQL (a part of the LAMP stack described in Tip Ten) was part of Oracle’s purchase of Sun Microsystems. Since Oracle sells its own, proprietary databases, MySQL developers were concerned about its future. To protect and continue to improve it, they forked the code into a variety of flavors and new projects.⁶⁵

A similar thing can happen to non-software, open-licensed projects. For example, in 2002 the Spanish language team of Wikipedia took a copy of Spanish Wikipedia content and created its own encyclopedia, *Enciclopedia Libre Universal*.⁶⁶

The implications are important for government and nonprofit managers, as a fork can represent loss of control—some part of the development community has objected and wants to go down another path. However, the existence of this threat gives contributors power and reassures them that they will always have access to their contributions. New forks also have the potential to spawn innovations that would not have come from one track alone.

Creative projects can provide more than just reusable source code and content. Using simple tools like RSS or more complicated application programming interfaces (APIs), projects can provide on-demand access to dynamic services. It is this approach that allows website developers to use and create appropriate context for maps using Google’s map interface.

62. Kevin Kelly. *Out of Control*. <http://www.kk.org/outofcontrol/ch24-a.html>

63. Christopher Kelty. *Two Bits: The Cultural Significance of Free Software*. 2008, p. 10.

64. Eric Raymond. “Homesteading the Noosphere.” August, 8, 2002, <http://catb.org/~esr/writings/homesteading/homesteading/ar01s09.html>

65. Nathan Willis. “A look at the MySQL forks.” *LWN.net*. April 22, 2009. <http://lwn.net/Articles/329626/>

66. Nathaniel Tkacz, “The Spanish Fork: Wikipedia’s ad-fueled mutiny.” *Wired*, January 20, 2011. <http://www.wired.co.uk/news/archive/2011-01/20/wikipedia-spanish-fork>

Tip Six: Reach for the Edges

Out on the edge you see all kinds of things you can't see from the center.

—Kurt Vonnegut⁶⁷

Request for Comment (RFC) 1958 was pretty specific. It said, “The network’s job is to transmit datagrams as efficiently and flexibly as possible. Everything else should be done at the fringes.”⁶⁸ Internet designers’ decision to support activities “at the fringes” is an important philosophic insight worth following. Open projects generally should seek for decisions and contributions to be made at remote edges of their systems.

Why are contributions from the edges important? With complicated problems, understanding of the problem and possible solutions are beyond our perception horizon—we need observations, responses, and insights from the edges to solve them.

Science author Steven Johnson quotes computer scientist Christopher Langton, who explained that “innovative systems have a tendency to gravitate toward the ‘edge of chaos:’ the fertile zone between too much order and too much anarchy.” At these edges are different kinds of people with different insights and experiences wrestling with different kinds of problems in new ways. John Hagel uses similar language to explain: “Edges are places that become fertile ground for innovation because they spawn significant new unmet needs and unexploited capabilities and attract people who are risk takers. Edges therefore become significant drivers of knowledge creation and economic growth, challenging and ultimately transforming traditional arrangements and approaches.”⁶⁹

There are a few reasons for the fertility Hagel describes. Participants not in the “core,” using Hagel’s word (if you are running a project, you are probably in the core), have different experiences and data, are more likely to be outside of current thinking, and may bring new insights. They are also likely to have different incentives and to be more interested in creating and using new flows of knowledge than preserving existing stocks.

Steven Johnson explains what edge participants are wrestling with, borrowing the term “the adjacent possible”⁷⁰ from biologist and Professor Stuart Kauffman of the University of Vermont. The adjacent possible describes the set of opportunities created because of ideas and tools available at the same time. Thus, where there is hydrogen and oxygen, water becomes an adjacent possible.

Johnson argues, “Good ideas are not conjured out of thin air; they are built out of a collection of existing parts, the composition of which expands (and, occasionally, contracts) over time.”⁷¹ He uses the discovery of oxygen as an example which required both existence of a set of ideas about air and gases as well as existence of advanced equipment—precise scales—able to measure minuscule changes in weight. When the ideas and the tools were both available, identifying oxygen became possible.

In order to respond to complicated problems, we need to increase the realm of the adjacent possible. To expand the realm of the adjacent possible, we want to engage the brains, resources, and insights of people at the edges of our problem space as they wrestle to bring together the knowledge and tools they are in contact with. Thus, we want to encourage participation by a

67. Kurt Vonnegut. *Player Piano*, 1952, p. 84. (ht: @socialledge)

68. “Architectural Principles of the Internet.” RFC 1958. June 1996. <http://www.ietf.org/rfc/rfc1958.txt>

69. Hagel, p. 16.

70. Steven Johnson. *Where Good Ideas Come From*. 2010, p. 31.

71. *Ibid*, p. 35.

wide variety of people in different settings (Tip One and Tip Nine) and build multiple communications channels as described in Tip Four. We want to engage and support edge participants and encourage them to provide fresh ideas and feedback.

Tip Seven: Make it Work, Then Make it Better

A complex system that works is invariably found to have evolved from a simple system that worked. The inverse proposition also appears to be true: A complex system designed from scratch never works and cannot be made to work. You have to start over, beginning with a working simple system.

—Gall's Law⁷²

While Gall's Law is written in the spirit of the Peter Principle,⁷³ with tongue in cheek, it contains more than a grain of wisdom. It was echoed in the now very familiar RFC 1958 as "Keep it simple. When in doubt during design, choose the simplest solution."⁷⁴ When trying to engage and coordinate large numbers of distributed people and groups, simplicity matters. It makes understanding easier, eases learning, and improves communications. Moreover, simpler things are easier to get running.

The Internet is in a constant state of iterative improvement. Examples of this evolution range from the development of Internet protocols from TCP to TCP & IP, or the multiple generations of e-mail or the evolution of newsgroups or the five changes an hour being made to Linux or the constant editing of Wikipedia. Likewise, efforts like Gopher, WAIS, and Pointcast have died away, replaced by new generations of services.

Mark Zuckerberg, Facebook's first software developer and majority owner, calls this the Hacker Way. In a public letter in preparation for Facebook's IPO, he says "The Hacker Way is an approach to building that involves continuous improvement and iteration. Hackers believe that something can always be better, and that nothing is ever complete. They just have to go fix it—often in the face of people who say it's impossible or are content with the status quo."⁷⁵

Making something work has been a consistent theme of Internet evolution. What we know as the Internet today is based on the TCP/IP networking protocol. This is what computers use to communicate with each other and pass data back and forth. It replaced the ARPANET's version, called the network control program (NCP), in 1983.

However, TCP/IP wasn't supposed to become the global networking standard. While it was being developed and implemented, the International Organization for Standardization (ISO) was developing a competing framework called OSI. In 1988 *Network World* ran the article, "Users must prep for TCP/IP-to-OSI move: Experts say migration to standard is inevitable"⁷⁶ and reported "while TCP/IP has a huge and growing base, it is bound to be uprooted and supplanted by OSI ... because OSI applications will offer more functionality, while the lower layers will be more reliable and easier to customize for any particular user's internetworking demands. In addition, more vendors will support OSI than TCP/IP..." That same year, the U.S. Government even announced that it was adopting this new protocol.

72. "Gall's Law." Wikipedia article. http://en.wikipedia.org/wiki/Gall%27s_law

73. "The Peter Principle." Wikipedia article. http://en.wikipedia.org/wiki/Peter_Principle

74. "Architectural Principles of the Internet." RFC 1958. June 1996. <http://www.ietf.org/rfc/rfc1958.txt>

75. Mark Zuckerberg. "Facebook's Letter from Mark Zuckerberg." *The Guardian*. February 1, 2012. <http://www.guardian.co.uk/technology/2012/feb/01/facebook-letter-mark-zuckerberg-text>

76. "Users must prep for TCP/IP-to-OSI move: Experts say migration to standard is inevitable." *Network World*, April 11, 1988 p. 15+. <http://books.google.com/books?id=IBMEAAAAMBAJ&lpg=PA15&ots=wXy144hjpP&dq=%22osi%20will%20replace%22&pg=PA15#v=onepage&q=%22osi%20will%20replace%22&f=false>

But it never happened. TCP/IP's widespread adoption persisted and continues to grow today. OSI was defined, but never widely implemented. Vint Cerf explains that "everything about OSI was described in a very abstract, academic way ... The language they used was turgid beyond belief. You couldn't read an OSI document if your life depended on it." He also reflected on the importance of having TCP/IP up and running. "We felt compelled to try things out, because in the end there was no point in specifying something if you weren't going to build it. We had this constant pragmatic feedback about whether things worked or didn't."⁷⁷

There is a similar, "working is better," story from Linux. Torvalds himself admitted, "From a theoretical (and aesthetical) standpoint Linux loses. If Richard Stallman's GNU kernel had been ready last spring, I'd not have bothered to even start my project: the fact is that it wasn't and still isn't. Linux wins heavily on points of being available now."⁷⁸

"Making it work" can play a practical role in projects that have open participation by forcing suggestions to be made real and critics to become contributors. Arien Malec warns that "many processes get bogged down by theoretical disagreements." To avoid this, open projects respond to suggestions, says Malec, by saying "OK, do it" and thereby separating the "lawyers from doers." He says this will encourage participation and increase learning "from the encounter with the real world." During the NwHIN Direct Project, Malec says, "we learned so much from four to six weeks of implementing [transport protocols]. We learned where the real hard problems actually were—not in transport but in trust."⁷⁹

With open systems development we are striving for what Aneesh Chopra, former White House Chief Technology Officer, calls focused collaboration, representing how government can work with other parties to innovate. He argues that the approach is more open and participatory than traditional command and control projects, but more mission and outcome-oriented than "let a thousand flowers bloom" approaches.

Keeping the focus simple is a key part of enabling participation (Tip Ten) that helps focused collaboration work.

Tip Eight: Make it Work, Then Standardize

We reject: kings, presidents and voting. We believe in: rough consensus and running code.

—David D. Clark⁸⁰

Coordinating the work of many people who may not know or talk directly with each other is a tricky problem. We do this all the time in a market economy, relying on prices to communicate and coordinate. Likewise, with language, we've agreed to the meaning of words so we can use them to communicate. The Internet relies on many agreements, called protocols, to allow computers and software to communicate and cooperate.

In fact, much of what we think of as the Internet is really just common language—standards and agreements—to send and receive data or behave in a certain way. The data may connect computers to one another or request and return a web page. The agreements which today make up the Internet include TCP/IP (Transmission Control Protocol/Internet Protocol), FTP

77. Hafner and Lyon, p. 247

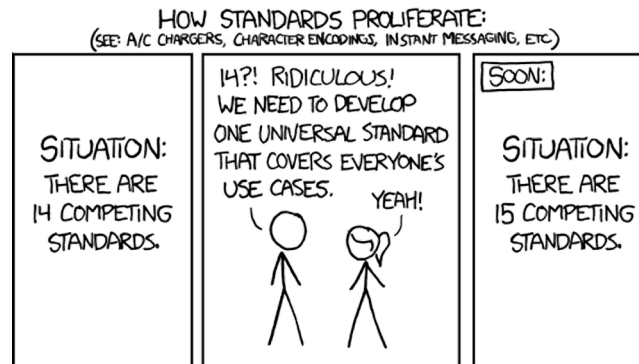
78. Ryan Paul. "March of the Penguin: Ars looks back at 20 years of Linux." *Ars Technica*. Aug. 25, 2011. <http://arstechnica.com/open-source/news/2011/08/march-of-the-penguin-ars-looks-back-at-20-years-of-linux.ars>

79. Author communication with Arien Malec.

80. "David D. Clark." Wikipedia article. http://en.wikipedia.org/wiki/David_D._Clark#cite_note-1 Original: slide 19 on PDF p. 551 of <http://www.ietf.org/old/2009/proceedings/prior29/IETF24.pdf>. July 1992.

(File Transfer Protocol), SMTP (Simple Mail Transfer Protocol), and POP (Post Office Protocol). Similarly, the World Wide Web relies on HTML (HyperText Markup Language) and HTTP (HyperText Transfer Protocol).

Figure 3: Standards proliferation⁸¹



Many of these protocols were created by the Network Working Group (NWG), whose main challenge was to agree in principle. The rules of the NWG in the initial stage were simple—“anyone could say anything and that nothing was official.” The fact that the NWG had no formal structure and as an entity didn’t really exist reinforced the message that standards could not be imposed. The role of a standards proposer was to explain value, demonstrate operations, and encourage adoption.

Demonstrating functioning implementation is an important criterion for getting agreement on and adoption of a new protocol. Moreover, the goal of these standards was not to strictly define processes, but to communicate and coordinate. The standards weren’t seen as an academic product but as a necessary and malleable component of a working system.

Jon Postel, a central member of the NWG, served as RFC editor from 1969 until he died in 1998. Among Postel’s many contributions to the Internet was Postel’s Law, which stated “Be conservative in what you send; be liberal in what you accept.” A little flexibility increases the chance that communications will succeed—the goal is to make it work.

Tip Nine: Take Advantage of All Types of Organizations

My name is Linus Torvalds and I am your god.

—Linus Torvalds⁸²

A byproduct of letting everyone play (Tip One) is the likely mix of organizations that may add value to collaborative projects. An interesting project is likely to require contributions from an ecosystem of government, nonprofit, business, academic, and individual participants with varying interests and incentives.

The Internet has been developed by a wonderful mishmash of people working in a wide variety of professional settings. It was started with U.S. Department of Defense funding through an Advanced Research Projects Agency (ARPA) project as a way for founding father Bob Taylor

81. Xkcd Standards. <http://xkcd.com/927/>

82. Linus Torvalds. “Jokingly introducing himself, at the 1998 Linux Expo in Durham, North Carolina,” http://en.wikiquote.org/wiki/Linus_Torvalds

to address his irritation at having multiple terminals in his office connected to different remote computers.⁸³ Taylor recruited Larry Roberts from MIT's Lincoln Labs to run the new program. A commercial consulting firm, BBN, was hired to develop the proof-of-concept, first generation ARPANET, relying heavily on insights from academics and labor from students at MIT and UCLA. Network thinkers from Stanford Research Institute, RAND, Britain's National Physical Laboratory, and elsewhere compared ideas and contributed. The original equipment for the first generation routers was manufactured by Honeywell.

A generation of for-profit enterprises has already bloomed and faded, including Netscape with its web browser (now evolved to the open source Firefox product); Pointcast, an early push news service; search engines like Alta Vista, Hotbot, and Excite; and a variety of retail attempts like Webvan, eToys, and Pets.Com. These companies have been replaced by a new generation of dominant Internet businesses including Google, Amazon, and Facebook, as well as many traditional bricks-and-mortar businesses expanding online.

The federal government no longer plays a central role in funding the Internet, though it is in continual discussions about how much control it should have. Of course, it is also implementing its own online services, increasingly opening up online access to data and materials, and regulating telecommunications and intellectual property. Other national governments, notably those of China and Iran, attempt to exert more direct control over what parts of the Internet are available to their residents.

The Internet is not just government and business. Individuals, informal groups, and nonprofits have been and continue to be instrumental in its operation and growth. Many groundbreaking ideas and products have come from amateur hackers, students, and individual entrepreneurs. A few of the many examples include versions of NNTP from Duke University students, IRC and Linux from Finnish university students, and HTML (and thus the World Wide Web) from a British physicist working at CERN.

Nonprofit and informal organizations play increasingly critical roles as custodians of valuable assets—software, patents, standards, brands, and more—managing ongoing development and investment, coordinating resources, and conducting negotiations. This includes groups like these:

- The IETF, which develops and promotes Internet standards (the modern version of the Network Working Group)
- The W3C, which is the main standards organization for the World Wide Web
- The Apache Software Foundation, which supports the Apache web server
- The Wikimedia Foundation, which coordinates and raises funding for Wikipedia
- The Linux Foundation, which supports the Linux operating system

Giving neutral public sector entities like foundations or nonprofits the responsibility for shared assets like software and standards can be crucial to their adoption. David Winer's company gave control of the RSS specification to Harvard's Law School so that competing companies would be more open to implementing it. The most recent version of RSS is released by Harvard under a Creative Commons license.

More than 20 years of development of Linux shows the potential for coordination between various flavors of people and organization. Though started by Torvalds when he was a student,

83. I relied on "Wizards and Internet Pioneers" (<http://www.ibiblio.org/pioneers/index.html>) for this recounting.

business has been interested in Linux from early on. While he was getting started, computer manufacturer DEC helped cover his costs to one of their user conferences and gave him use of one of their new computers. Red Hat, the first distributor of Linux and now a publicly traded company, offered its first distribution in 1994.⁸⁴

Actual development of the Linux software shows the same pattern of cooperation across sectors and erstwhile competitors. According to its foundation, the huge Linux code base (over 14 million lines of code) has contributions from 500 separate companies. Nonetheless, in the last accounting, almost one in five contributions come from independent developers, more than any individual company. Red Hat was the most active business contributor, accounting for 12 percent of changes, and was followed by Intel, Novell, and IBM.⁸⁵

Zittrain emphasizes the importance of participation by many kinds of organizations and individuals, explaining that “generatively enabled activity by amateurs can lead to results that would not have been produced in a firm-mediated market world.”⁸⁶ Benkler goes even farther, saying “We are seeing early signs of the emergence of an innovation ecosystem made of public funding, traditional nonprofits, and the newly emerging sector of peer production that is making it possible to advance human development through cooperative efforts ...”⁸⁷

Tip Ten: Design for Participation

Modularity is good. If you can keep things separate, do so.

—RFC 1958⁸⁸

Coordinating activities by multiple contributors requires developing a project structure that lets individuals and small teams contribute constructively, even when they are new to a project. Tim O’Reilly calls this creating an “architecture of participation.”⁸⁹ A number of approaches make it easier for others to contribute including open communications (see Tip Three and Tip Four), sharing source code (Tip Five) and creating a sharp focus and simple objective (Tip Seven). Careful design of the project’s structure can also enable more contributions—break the work into manageable pieces that can be stacked and create multiple roles for contributors.

Breaking work into manageable pieces makes it easier for individuals or small teams to understand what needs to be done, spread their labor across activities, and climb the learning curve to become involved. Benkler defined two dimensions that support layering and distributed contributions—*modularity* and *granularity*. Modularity is the ability to break a project down into separate pieces and *granularity* refers to the size of the pieces. Smaller granules are, in general, easier to contribute and invite more participation though they may be harder to aggregate into usable modules. So the trick is to define granules that are small and easy to create frequently, while still being simple and manageable enough to aggregate into valuable modules.⁹⁰

Software programmers are trained to develop software in modules. Using modules helps the programmer break down complicated problems into a set of more manageable pieces.

84. Ryan Paul. “March of the Penguin: Ars looks back at 20 years of Linux.” *Ars Technica*. Aug. 25, 2011. <http://arstechnica.com/open-source/news/2011/08/march-of-the-penguin-ars-looks-back-at-20-years-of-linux.ars/2>

85. Ryan Paul. “Linux Kernel.” *Ars Technica*. December 1, 2010. <http://arstechnica.com/open-source/news/2010/12/linux-kernel-13-million-lines-over-5-patches-per-hour.ars>

86. Zittrain, p. 84.

87. Yochai Benkler, *The Wealth of Networks*. Yale University Press, 2006. p. 15

88. Network Working Group, “Architectural Principles of the Internet.” Request for Comment 1958, June 1996. <http://www.ietf.org/rfc/rfc1958.txt>

89. Tim O’Reilly. “The Architecture of Participation.” June 2004. http://oreilly.com/pub/a/oreilly/tim/articles/architecture_of_participation.html

90. Benkler, pp. 100–101.

Modules share their functions while hiding their complexity, allowing other programmers to treat them as black boxes—they work and we don't need to know how. Modules can improve thinking, encourage reuse of code, and improve a program's readability and maintainability.

Similarly, groups of programs can work together and support projects in what is commonly called a stack. Stacks are an important reason why the Internet has managed to grow and evolve so quickly. Stacks are made up of a set of protocols, software, and hardware that, when joined together, provide a particular service or set of services.

The Internet is full of stacks, starting with the Internet Protocol Stack that defines how data passes through the network from an application on one computer to an application on another computer. Another influential example is the LAMP Stack, where LAMP describes a set of open source software products—the Linux operating system, Apache web server, MySQL database software, and a programming language (Perl or PHP or Python). While the software products have been developed separately, the set of them used together makes a powerful, low-cost set of tools for developing websites.

The Internet Protocol Stack	
Layer	Sample Protocol
Application	SMTP
Transport	TCP
Network	IP
Link	Wi-Fi
Physical	10 Base T

It isn't only software that uses stacks—they are woven throughout our culture. Johnson uses the example of a stack of fields of science writing: “The field that ultimately explained Darwin's Paradox—ecosystems ecology—stands on the shoulders of population genetics, systems theory, and biochemistry, among others.”⁹¹

It is this potential—to have separate teams solving separate problems at the same time and then stitch together the results to solve even bigger problems—that makes the stack so compelling. With a stack structure, program managers can hope to coordinate large numbers of contributors—many groups working simultaneously on their own.

Another way to encourage contributors is to allow multiple kinds of contributions and let people contribute in multiple ways. Hagel recommends that, rather than specify activities for teams, the focus should be on the outcome of their process. He explains this provides “more space for individual participants to experiment, improvise, and innovate.”⁹² By not specifying specific tasks or roles, we allow what Shirky calls “spontaneous division of labor.”⁹³

Wikipedia contributors can choose many levels of involvement. They can dip into a page and fix spelling or a comma, specialize in creating links to other articles, be part of the grammar police who improve wording in articles, monitor articles for vandalism, take on responsibility for drafting new articles, or become editors providing guidance and support to other contributors.

91. Johnson, p. 191.

92. Hagel, p. 196.

93. Shirky, p. 118.

Red Hat's Fedora Project, which leads development of a home and small business version of Linux, is explicit about some of the roles it is looking for from contributors. It calls for people to be content writers, designers, operating system developers, translators, web developers and administrators, or "people people" involved in marketing and outreach.⁹⁴

Allowing flexibility in contributor roles encourages people to offer their specific skills and interests to the project.


Tip Eleven: Increase Network Impact

Mathematics says the sum value of a network increases as the square of the number of members.

—Kevin Kelly⁹⁵

By many metrics, Internet growth is stunning. Whether counting number of computers connected, people using, websites hosted, e-mails sent, or advertising dollars spent (see Figure 4), growth has been on an upward sloping curve. How can nonprofit and government leaders generate similar growth in value from their projects? One approach is to deliberately look for ways to increase network impacts.

Figure 4: Quarterly Growth in Internet Advertising 1999–2011




Source: IAB⁹⁶

Network impacts derive from the ability of network participants to affect many other network participants. We can visualize three types of impact based on the type of interconnections supported (Figure 5). The first, traditional model is the one-to-many network broadcast communications model. This model, used by TV and radio stations and adopted by newsletter writers on the Internet, attracts audiences, subscribers, or members and sends them information directly. The audience grows member by member and the effort required to get a member generates one member's worth of value. We can think of these approaches as approximating linear growth in value.

94. "Join Fedora." <http://fedoraproject.org/wiki/Join>


95. Kevin Kelly. *New Rules for the New Economy*. 1999. See <http://www.kk.org/newrules/newrules-2.html>

96. IAB Internet Advertising Revenue Report. April 18, 2012. http://www.iab.net/media/file/PwC_Presentation_FY2011_print.pdf slide 7

Figure 5: Three Types of Network Structure

A second community model is a network of nodes that can interact back and forth with each other. Arun Sundararajan, professor at NYU's Stern School of Business, explains, "A product displays positive network effects when more usage of the product by any user increases the product's value for other users."⁹⁷ The telephone is a simple example. One phone on a network is not very useful, but two phones start to get useful and participants get more value as more phones are added. It is important to reinforce that the additional value isn't just to the new phone owner. Previous owners get a bit more value as well. It is a kind of self-reinforcing feedback loop—as the number of users grows, the system becomes more valuable, so the number of users grows.⁹⁸

This phenomenon tends toward an exponential growth curve like that shown in Figure 6. Engineer and inventor Robert Metcalfe used a network effects argument to sell his Ethernet networking hardware, saying that the value of his hardware increased with the square of the number of cards installed. This relationship is known as Metcalfe's Law. John Hagel uses Microsoft, FedEx, the Internet, and fax machines as examples of products that display network effects.⁹⁹

Figure 6: Metcalfe's Law—Network Value Increases Faster than the Size of the Network

97. Arun Sundararajan, New York University. Blog. <http://oz.stern.nyu.edu/io/network.html>

98. The system may stop growing or even decline beyond the point where new users bring zero or negative additional value.

99. John Hagel III and Arthur G. Armstrong. *Net Gain*. 1997, p. 43.

Network innovator David Reed suggests a third type. He argues that even Metcalfe's Law underestimates the possibility of network effects, especially in cases where groups can define their own sub-networks. He calls these group forming networks, or GFN, suggesting a GFN allows "small or large groups of network users to coalesce and to organize their communications around a common interest, issue, or goal."¹⁰⁰

These three types of structure have been represented in various stages of Internet development. One example is in the connectivity of networks themselves as we've evolved from connections to one computer (type 1), to connecting computers to each other on one network like ARPANET (type 2), to connecting networks to other networks via TCP/IP (type 3). This argument helps explain the power of Internet growth as TCP/IP was adopted and what had been a set of discrete networks, including ARPANET, became an interconnected network-of-networks.

There is ongoing discussion of how to actually measure network effects, negative impacts, and limits to growth.¹⁰¹ However, the basic insight is that increasing connectivity, connecting more groups, and increasing flexibility in defining groups will tend to increase value.

Tip Twelve: Build Platforms

The [interstate highway system] was a triumph of platform thinking, a key investment in facilities that had a huge economic and social multiplier effect. Though government builds the network of roads that tie our cities together, it does not operate the factories, farms, and businesses that use that network...

—Tim O'Reilly¹⁰²

Platforms are tools that help people and organizations coordinate their activities so they are jointly more productive. Hagel defines platforms as "frameworks for orchestrating a set of resources that can be configured quickly and easily to serve a broad range of needs."¹⁰³ Brian Behlendorf describes platforms as plumbing that provides innovators access to a set of services they need in order to innovate, and usually would prefer not to build on their own.¹⁰⁴

Successful platforms have a multiplier effect—the value they create comes from other projects. In the social space, a platform project is not trying to solve all problems, but to provide resources to help other people solve them. Platforms allow others to build more powerful solutions and services more easily.

Many of the platforms we're familiar with are based on software and computer hardware, though the potential of platforms is not limited to software development. While Microsoft with Windows, Apple with its App store, and Linux have created powerful platforms, so has the neighborhood shopping mall. An important element of platforms is that they help different groups connect to create value. Thus, Apple connects software developers and iPhone users through its Apps store the same way shopping malls, by providing parking and storefronts, help retailers and their customers connect.¹⁰⁵ Arguably, the market is the world's biggest

100. <http://www.reed.com/dpr/?sel=gfn/>

101. See Krugman's discussion of whether "increasing returns" are inevitable (<http://web.mit.edu/krugman/www/metcalfe.htm>) as well as Beckstrom's extensions to Metcalfe's law suggesting potential negative effects (http://www.beckstrom.com/Beckstrom%27s_Law) as well as David Reed's argument that Metcalfe may underestimate potential impacts (<http://www.reed.com/dpr/locus/gfn/reedslaw.html>)

102. Tim O' Reilly. "Government as a Platform." *O'Reilly Open Feedback Publishing System*. http://ofps.oreilly.com/titles/9780596804350/defining_government_2_0_lessons_learned_.html

103. Hagel, p. 75

104. David Witzel. "EDFux Call #8 Afterthoughts." *EDF + Business*. Blog. <http://blogs.edf.org/innovation/2010/04/11/edfix-call-8-after-thoughts-building-plumbing-for-business-sustainability/>

105. David S. Evans, Andrei Hagiu, Richard Schmalensee. *Invisible Engines*, 2008, p. 59

platform. Its simple price structure allows buyers and sellers to coordinate and cooperate even when they don't know each other.

Windows software and Microsoft's partnership with Intel created a powerful platform bringing together four groups of value creators. The Win/Tel partnership supported many different computer vendors including Compaq, IBM, and Dell; software developers writing applications for businesses and homes; peripheral hardware developers selling printers, disk drives, modems and more; all of which supported the fourth group—computer users.¹⁰⁶

The Web is a massive platform built on open standards. As O'Reilly points out: "Tim Berners-Lee didn't develop hundreds of millions of websites."¹⁰⁷ He designed a platform so that others could.

Tip Ten focuses on how to structure what we're building, keeping it modular and stepwise to allow participation, learning, and evolution. By imagining platforms, we focus on providing services to other problem-solvers—raising the floor so other efforts can be more successful, more quickly, at lower cost.

Platform thinking suggests a number of different ways to create this multiplier effect, including:

- Defining public standards to simplify communications and coordination between projects (Tip Eight)
- Creating reusable, open source products like Linux, which can be reused by other projects (Tip Five)
- Providing programming interfaces (APIs) that allow direct access to data and services (Tip Five)

One way to understand platform development is that the goal isn't necessarily to solve a problem but to give others the resources to help solve their problems.

106. Evans, p. 95

107. Tim O' Reilly. "Government as a Platform." *O'Reilly Open Feedback Publishing System*. http://ofps.oreilly.com/titles/9780596804350/defining_government_2_0_lessons_learned_.html

Case Study: Open Project Design at the Office of the National Coordinator for Health Information Technology

The Department of Health and Human Service's Office of the National Coordinator for Health IT (ONC)¹⁰⁸ is working to expand adoption of electronic health records, connect physicians to each other and to a national information network, involve patients in their health care, and broadly improve public health.¹⁰⁹ To achieve these ambitious goals, ONC has been using Internet-like, open project approaches in a number of programs and even outlined similar concepts in their Design Principles document that has guided recent project efforts.¹¹⁰ This case study provides a brief overview of two initiatives, using ONC's experience to illustrate what the tips from Internet pioneers look like in practice.

Background

In 2004, President George W. Bush called for widespread adoption of interoperable electronic health records and established the Office of the National Coordinator for Health IT to make it happen. The mission and scope of ONC greatly expanded in 2009 with the passage of the American Recovery and Reinvestment Act (ARRA), which established the Medicare and Medicaid electronic health record (EHR) incentive programs providing billions of dollars to providers and hospitals that adopt and "meaningfully use" health IT. The first set of requirements for meaningful use include recording medications, sharing patient health records for referrals, implementing clinical decision support, and reporting quality results.¹¹¹

ONC has worked with collaborators to define standards, services, and policies for communicating health information, collectively dubbed the Nationwide Health Information Network (NwHIN). ONC's goal with NwHIN is to "provide a secure, nationwide, interoperable health information infrastructure that will connect providers, consumers, and others involved in supporting health and healthcare. This critical part of the national health IT agenda will enable health information to follow the consumer, be available for clinical decision-making, and support appropriate use of healthcare information beyond direct patient care so as to improve health."¹¹² What's at stake? Better care for patients, reduction in medical errors and duplicate testing, more accurate diagnoses, and less time spent on administrative tasks for providers and patients alike.

108. Disclaimer: my wife, Claudia Williams, works in ONC and this case study is informed by conversations with and guidance from her.

109. Summarized from HHS handout, "Office of the National Coordinator for Health Information Technology (ONC)." <http://www.himss.org/handouts/executiveSummary.pdf> and Executive Order 13335. http://nodis3.gsfc.nasa.gov/displayEO.cfm?id=EO_13335_

110. "Design Principles." The Direct Project Wikipedia article. <http://wiki.directproject.org/Design+Principles>

111. "EHR Incentive Programs." Website of the Centers for Medicare and Medicaid Services. https://www.cms.gov/ehrincentiveprograms/30_Meaningful_Use.asp

112. "Nationwide Health Information Network (NHIN): Background & Scope." HHS website. <http://www.hhs.gov/healthit/healthnetwork/background/>

ONC's challenge is to get a wide variety of organizations including more than two dozen federal agencies, state and local governments, hospitals, insurers, patients, and some 800,000 doctors, to coordinate so that:

- Health data are securely shared
- Health care and health improve
- Costs go down

Major components of the NwHIN effort include:

- Negotiation of definitions, standards, and policies to support various kinds of information exchange
- Development, prototyping, and adoption of software to support exchange engagement with a diverse community of health professionals, vendors, patients, and government staff to refine, test, and implement these approaches

Accomplishments by NwHIN to date include:

- A community of federal agencies and private organizations working to develop and implement information exchange
- An open source tool to automate health quality reporting, called PopHealth
- A set of protocols for health information exchange, called NwHIN Exchange
- An open source software package that implements these protocols called CONNECT
- Simple, e-mail-based standards for sharing health data called Direct

The following cases focus on the last two of the activities listed above:

- CONNECT Software: a free, open source software solution that supports health information exchange
- Direct Project: standards and services to enable secure health information exchange

CONNECT Software

CONNECT is the current incarnation of a software development effort that started in 2007 to share health-related data among the more than 20 federal agencies that house and use it. It is an open source package that implements NwHIN standards. The project claims “more than 2,000 organizations—including federal agencies, states, healthcare providers, insurers, health IT vendors—all working together to improve the CONNECT solution.”¹¹³

Started as a traditional government software effort developed by interested federal agencies under the auspices of the Federal Health Architecture, CONNECT was developed under a contract with Harris Corporation in 2008. The software was released as open source code in 2009 under a license that places few restrictions on modification or redistribution. As of March 2012, it is in release version 3.3. The Department of Veterans Affairs, Department of Defense, and Thayer County, Nebraska are among the organizations that have begun to implement components of the CONNECT package.¹¹⁴

113. “The CONNECT Open Source Solution: A Gateway to the Nationwide Health Information Network.” http://healthit.hhs.gov/portal/server.pt/gateway/PTARGS_0_12811_953718_0_0_18/CONNECT_Data_Sheet.pdf

114. “Healthbridge Rolls Out Meaningful Use Exchange.” Mirth Corporation, November 10, 2009. <http://www.mirthcorp.com/news/healthbridge-meaningful-use-with-mirth>

The CONNECT project was actively engaged in building a software developer community for its product, including presenting at OSCON, a major conference for open source development, and hosting a series of code-a-thons around the country. In addition to providing access to all source code online,¹¹⁵ the project offers a wiki,¹¹⁶ online forums,¹¹⁷ and online issue tracking¹¹⁸ open to contributors. Project managers are available for direct contact and questions, the projects host webinars and training seminars, and even architecture discussions are held on open conference calls.¹¹⁹

In May 2011 a forked version of the software was released as Aurion in a project managed by the newly formed Alembic Foundation.¹²⁰ Started by two of the original CONNECT architects, the new foundation is developing open source software starting from a CONNECT codebase while extending decision-making for the direction of development beyond federal agencies.

ONC has announced a new generation of Connect, Connect+, that will, among other things, be restructured to simplify deployment.¹²¹ In 2012 ONC intends to spin off both NwHIN Exchange and the Connect software into external, public-private organizations. These new homes will manage governance and planning as well as development of the software.¹²²

NwHIN Direct

As the *New York Times* explains, “A basic challenge is for doctors, hospitals, patients and public health authorities to be able to easily and securely share information—things like a person’s vital signs, diagnosis, lab tests and drugs prescribed. A fancy electronic patient record, unconnected, is just an expensive way to capture data.”¹²³

The NwHIN Direct Project,¹²⁴ launched in 2010, is intended to address this challenge and replace the fax machine with an easy, standardized way to exchange health care information electronically. Currently, even doctors using the same brand of health record software might not be able to exchange information, while those with different brands of software have little hope of electronic exchange. Direct is intended to improve quality, speed transfer, and security, and lower costs for sharing patient data.

With Direct, ONC took an even more dramatic step toward open systems development. Foregoing responsibility for any software development, ONC focused on defining technical standards and services with the expectation that project partners would implement.

Representatives from electronic health records vendors, medical organizations, health delivery networks, federal and state governments, and consultants have participated.¹²⁵ Work has

115. “Source Code Repository.” CONNECT Community Portal. <https://developer.connectopensource.org/display/CONNECTWIKI/Source+Code+Repository>

116. Connect Community Wiki. <https://developer.connectopensource.org/display/CONNECTWIKI/Wiki>

117. “Developer Resources.” CONNECT Community portal. <http://www.connectopensource.org/developer-resources/forums>

118. <https://issues.connectopensource.org:8443/secure/Dashboard.jspa>

119. This video of NwHIN project managers presenting to Open Source developers at an open source conference, provides a flavor: <http://blip.tv/oreilly-open-source-convention/oscon-2010-arien-malec-NwHIN-direct-an-open-government-health-it-collaboration-3995791>

120. Aurion website. <http://aurionproject.org/>

121. <http://www.healthcareitnews.com/news/onc-releases-connect-33-looks-ahead-connectplus>

122. <http://www.healthcareitnews.com/news/nwhin-exchange-set-stand-its-own-oct>

123. Steve Lohr, *New York Times*. “U.S. Tries Open-Source Model for Health Data Systems.” February 2, 2011, <http://bits.blogs.nytimes.com/2011/02/02/u-s-tries-open-source-model-for-health-data-systems/>

124. Office of the National Coordinator for Health Technology. “Direct Project.” http://healthit.hhs.gov/portal/server.pt/community/healthit_hhs_gov_direct_project/3338

125. The Direct Project Overview, October 11, 2010. <http://wiki.directproject.org/file/view/DirectProjectOverview.pdf>

proceeded in a series of work groups convened by ONC, with over 200 organizations committed to active involvement and implementation.¹²⁶

Direct is based on widely adopted Internet standards, including e-mail protocols for transport (SMTP), content sharing (MIME), and security (X.509). The approach is intended to be widely applicable and able to be integrated with existing electronic health record packages and health information exchange platforms, as well as foster new products and services.

ONC had a clear focus to drive its efforts—coming up with an easily adopted mechanism for two providers to securely share patient health information. Arien Malec, former coordinator for NwHIN Direct, says the project has had an “iron mission, flexible tactics.” He explains, “We had a goal (universal addressing and transport for content-neutral, secure, directed exchange for health care) and stuck with it, through thick or thin, but were willing to try almost anything and compromise on almost anything to get there.”

This meant deliberately not including other potentially important but potentially distracting issues. The design principles of Direct Project state, “Don’t let ‘perfect’ be the enemy of ‘good enough;’ go for the 80% that everyone can agree on ... before focusing on the more obscure” interpreted as “when in doubt, cut it out.”¹²⁷ For example, questions about how to find doctors’ Direct addresses (similar to e-mail addresses), while interesting and probably valuable, was an issue deliberately defined as out-of-scope.

ONC has encouraged participation by a wide range of affected and interested organizations and people. Technology providers like Microsoft and Google, health data companies like Surescripts and Allscripts, representatives of medical providers like the American Academy of Family Physicians, consultants, and regional health agencies are involved.

To help manage participation, ONC has used a variety of communications channels. Arien Malec says they “pushed for ideas to be explored in working code and for talk to get written down.” They have used a blog¹²⁸ and more recently Twitter¹²⁹ to broadcast announcements to the community.

However, it is through their wiki¹³⁰ and code archives that much of the detailed learning is shared and practical decisions made. The Direct Project wiki, actively used for documentation and coordination, lists over 800 members. ONC has conducted face-to-face boot camps attended by state partners. Notes and recordings from those sessions are available online.¹³¹

One of the benefits of open participation has been innovation from the edges. Malec recalls being approached at a conference by someone wanting to participate in The Direct Project. Malec says, “He was from an obscure company in Minnesota that I’d never hear of before and clearly wasn’t on my hit list of folks I wanted to get involved. I told him: we’ve got a wiki; the best way to get involved is to go participate and add value.” The new participant proceeded to provide a helpful test model for how the architectural components could be provided and his company ended up being the first to launch in production.

126. Dr. Farzad Mostashari. “12 Months of Health Information Technology: A Year of Momentous Progress.” January 6, 2012. Health IT Buzz Blog. Office of the National Coordinator for Health Technology. <http://www.healthit.gov/buzz-blog/from-the-onc-desk/12-months-health-information-technology-year-momentous-progress/>

127. “Design Principles.” Direct Project wiki. <http://wiki.directproject.org/Design+Principles>

128. Health IT Buzz. Blog. Office of the National Coordinator for Health Technology. <http://www.healthit.gov/buzz-blog/>

129. http://twitter.com/#!/Farzad_ONC

130. The Direct Project Wiki: <http://wiki.directproject.org/> and the CONNECT wiki: <https://developer.connectopensource.org/display/CONNECTWIKI/Wiki>

131. “ONC Direct Boot Camp.” Direct Project wiki. <http://wiki.directproject.org/ONC+Direct+Boot+Camp>

NwHIN Direct operated as a “do-ocracy” to handle disagreements—if you’ve got a problem you’ve got to suggest a solution. Malec explains, “We had a unanimous, consensus-based approach where one no vote would stop discussion, *but* we demanded that no votes be accompanied with the fix to turn the no to a yes.”

ONC has not commissioned or paid for any software to be developed as part of this project. Instead, participants agreed to provide their implementations under open source licenses and have thus generated versions that run in various open and proprietary frameworks.

This code has been developed iteratively along with the specifications, letting participants identify and refine requirements based on actual implementation experience. Malec says their philosophy is to “ship the 1.0”—the first full version of code—because “you learn more by getting stuff out there and seeing what actually happens than you do by talking through what might happen.” He says “the specifications were then based on observed need, not anticipation of needs.”

Malec gives an example of the kind of learning they gained with the iterative process and early code release. He explained that a major focus, during the planning process, was on what protocol to use to move the data (e-mail vs. web vs. programming interfaces). But, after getting a version released, they were surprised because “Turns out, scalable, interoperable *trust* was the hard thing and the pros/cons of any of the transports was dwarfed by that issue.”¹³²

NwHIN Direct creates a new layer of the stack for health information exchange. This layer has already been used in many ways, including for public health reporting, sharing information when patients are referred, sending patients their health information, and sending secure messages within a hospital.

ONC Results: A Preliminary Assessment

While ONC’s work is still young, some early indicators are positive. The open source CONNECT Project reports that over 2,000 organizations have downloaded or contributed to it. Whether the forking of the CONNECT codebase to Aurion will advance HHS’s objectives is yet to be seen. It is possible that moving Exchange and Connect to private organizations will respond to external demands.

Meanwhile, live implementations of the Exchange protocols are in place across the nation, with 35 companies implementing in 2011 and over 40 states building Direct into their strategies to support health information exchange. Over two dozen companies are developing products and services to integrate or build upon ONC’s efforts.¹³³ Direct has been incorporated into the proposed regulations for the second stage of “meaningful use” to provide a ubiquitous mechanism for information exchange.¹³⁴

HHS’s Office of the National Coordinator is an example of a government agency taking the lead to explore, demonstrate, and benefit from open project approaches like those that built the Internet. Projects like CONNECT and Direct that involve diverse organizations, coordinate the efforts of many people, and create products that can be widely used and reused, demonstrate new ways of thinking about government service provision.

132. Conversation with the author.

133. Ability website. <http://www.abilitynetwork.com/services/Direct/>

134. <http://www.healthit.gov/providers-professionals/meaningful-use-stage-2>

Conclusion

Why did it take 5,000 years to put wheels on luggage?

—Author Unknown

Over the past 40 years, innovations during the development of the Internet have challenged and toppled conventional wisdoms—that we can't organize large informal groups, that people won't contribute for free, that enterprise software can't be built this way, that the Internet itself will be replaced or collapse, and more. Meanwhile, the practical, let-everyone-play and count on “running code” focus of the Internet's pioneers has led to approaches and products with amazing reach and influence. These insights have the potential to make dramatic contributions to how we can do what we need to do.

Development of the Internet has been project-based and iterative. A host of often volunteer-run projects has built, improved on, and replaced infrastructure and services; and tested, discarded, or institutionalized innovations. These resources tend to be reusable, providing infrastructure for more experimentation and progress in rapid cycles of creation.

The ability of these open project approaches to be successful comes from three broad multipliers of value:

- First, open projects are able to overcome resources constraints by harnessing resources—labor, money, equipment, or ideas—from many locales.
- Second, they have the potential to scale with network speed building upon not just networks, but networks of networks.
- Third, they foster innovation by encouraging participation, lowering the cost of experimentation, increasing the sharing of lessons (and failure) and reaching toward the edges for new data and possibilities.

Open approaches may be particularly valuable strategies for grappling with complex, “wicked” problems which are dynamic and for which no clear solution is known. They operate better in dynamic environments and are able to generate more innovation and experimentation. The ability to aggregate distributed efforts may make them better able to reach critical mass and scale.

Acknowledgements

I'd like to explicitly thank some of the many people who have helped me understand and try to explain the wonder that is the Internet and the exciting lessons that derive from its development. In particular, I'd like to thank:

- My former colleagues at Environmental Defense Fund; especially my boss, Gwen Ruta, who gave me space to work on these ideas; Justin Olsson, who was a supporter and contributor to this text; and Beth Trask, who helped me work through many of these thoughts in practice.
- Jane Fountain, professor and ongoing adviser.
- Jerry Michalski, fellow REX'rs and the generous participants in Jerry's list including David Gray, Peter Kaminski, Kevin Jones, Isabel Draves, Bill Seitz, Cameron Barrett, Mark Finnern, Richard Gabriel, Tom Mandel, and David Isenberg.
- My colleagues and partners at Forum One, especially Jim Cashel, who not only reviewed and commented but nudged me to finish.
- Arien Malec, RelayHealth, previously with ONC, who was generous with his time, experience, and ideas.
- The team at the IBM Center for The Business of Government, especially Mark Abramson and John Kamensky, who were constructive, patient, and encouraging.
- My wife, Claudia Williams, who has a passion for the subject, co-developed many of these concepts, introduced me to resources and people, inspires me to think harder, and sets high bars for accomplishment with her example.
- My boys Adam and Zach, who I hope will soon be contributing to the social stack.

Of course, all mistakes, misunderstandings, and misrepresentations remain solely my responsibility.

About the Author


David T. Witzel, BS, MPP, is an entrepreneur and organizational strategist. He is a Fellow with The EdgeLab, an ecosystem design firm based in San Francisco, California, and Chief of Conspiracy for the Green Innovators in Business Network, an organization he helped create to foster sustainable business change-makers.

He cofounded Forum One Communications Corporation, a web development firm based in Alexandria, Virginia, where he worked for 13 years. From 2009 to 2012, he was Director of the Innovation Exchange for the Environmental Defense Fund, where he explored new ways to engage businesses in environmental sustainability, working with organizations like InnoCentive, GE, and Ashoka.

Witzel has degrees in sociology from Texas A&M University and in public policy from Harvard's Kennedy School of Government. Most of his 30-year professional career has been at the intersection of computer technology, public policy, and organizational change. He learned to program using punch-cards and was a member of the Championship Texas A&M punch-card-stacking team in 1982.

In his first professional position, with the Texas State Data Center, he was responsible for processing data from the 1980 Census using IBM mainframes to read magnetic tape reels. He had a 300 baud acoustic-coupler modem at his house to support dial-up connections. He worked for start-up computer companies in New Jersey and Massachusetts on a variety of mini-computers in UNIX environments.

While in the Peace Corps in Botswana in 1987, he wrote and installed the first PC-based database to track teacher placements in the Ministry of Education. Later, while working for the Harvard Institute for International Development in Jakarta, Indonesia, he led installation of one of the first local area networks in the Ministry of Finance. He was a member of CompuServe and Gopher-user to get access to the newly public Internet and in 1993 implemented FidoNet in Jakarta to connect to store-and-forward e-mail.



He installed and managed one of the first web servers at the World Bank in 1995 while working in the environment section of the Research Department, and designed and launched the “New Ideas in Pollution Regulation” website demonstrating the web as a tool for transparency, data sharing, and knowledge dissemination. While at Forum One he continued work with the World Bank, leading efforts to launch the first database enabled websites, a members-only site for journalists, and *B-SPAN*, for sharing video of World Bank events. He also designed and led impact analysis of internal and external web activity.

At Forum One he led business and product development, leading efforts to launch the *Forum One Index*, a search engine for web discussions; *ProjectSpaces*, a tool to support team collaboration; *Live Interviews Online*, to host moderated question and answer events, and *Backyard Jungle*, a kid-focused website for PBSKids. He was a regular contributor to the *Online Community Report*.

He is a founding member of the *Falls Church Times*, the online community newspaper in Falls Church, Virginia, where he lives with his wife Claudia Williams and their two boys Adam and Zachary.

Key Contact Information

To contact the author:

David Witzel

305 Buxton Rd
Falls Church, VA 22046
571-641-3029

e-mail: dwitzel@policycommons.org
<http://www.twitter.com/dwitzel>



Reports from **IBM Center for The Business of Government**

For a full listing of IBM Center publications, visit the Center's website at www.businessofgovernment.org.

Recent reports available on the website include:

Assessing the Recovery Act

Key Actions That Contribute to Successful Program Implementation: Lessons from the Recovery Act by Richard Callahan, Sandra O. Archibald, Kay A. Sterner, and H. Brinton Milward

Managing Recovery: An Insider's View by G. Edward DeSeve

Virginia's Implementation of the American Recovery and Reinvestment Act: Forging a New Intergovernmental Partnership by Anne Khademian and Sang Choi

Collaborating Across Boundaries

Designing Collaborative Networks: Lessons Learned from Public Safety by Jane Fedorowicz and Steve Sawyer

Environmental Collaboration: Lessons Learned About Cross-Boundary Collaborations by Kathryn Bryk Friedman and Kathryn A. Foster

Conserving Energy and the Environment

Implementing Sustainability in Federal Agencies: An Early Assessment of President Obama's Executive Order 13514 by Daniel J. Fiorino

Breaking New Ground: Promoting Environmental and Energy Programs in Local Government by James H. Svava, Anna Read, and Evelina Moulder

Fostering Transparency and Democracy

Assessing Public Participation in an Open Government Era: A Review of Federal Agency Plans by Carolyn J. Lukensmeyer, Joe Goldman, and David Stern

Using Geographic Information Systems to Increase Citizen Engagement by Sukumar Ganapati

Improving Performance

Improving Government Contracting: Lessons from Bid Protests of Department of Defense Source Selections by Steven M. Maser

A Guide to Data-Driven Performance Reviews by Harry Hatry and Elizabeth Davies

A Leader's Guide to Transformation: Developing a Playbook for Successful Change Initiatives by Robert A. F. Reisner

Managing Finances

Strategies to Cut Costs and Improve Performance by Charles L. Prow, Debra Cammer Hines, and Daniel B. Prieto

Strengthening Cybersecurity

A Best Practices Guide for Mitigating Risk in the Use of Social Media by Alan Oxley

A Best Practices Guide to Information Security by Clay Posey, Tom L. Roberts, and James F. Courtney

Transforming the Workforce

Engaging a Multi-Generational Workforce: Practical Advice for Government Managers by Susan Hannam and Bonni Yordi

Implementing Telework: Lessons Learned from Four Federal Agencies by Scott P. Overmyer


Using Technology

Working the Network: A Manager's Guide for Using Twitter in Government by Ines Mergel

Reverse Auctioning: Saving Money and Increasing Transparency by David C. Wyld

Using Online Tools to Engage—and be Engaged by—The Public by Matt Leighninger

An Open Government Implementation Model: Moving to Increased Public Engagement by Gwanhoo Lee and Young Hoon Kwak



IBM Center for
The Business of Government

About the IBM Center for The Business of Government

Through research stipends and events, the IBM Center for The Business of Government stimulates research and facilitates discussion of new approaches to improving the effectiveness of government at the federal, state, local, and international levels.

About IBM Global Business Services

With consultants and professional staff in more than 160 countries globally, IBM Global Business Services is the world's largest consulting services organization. IBM Global Business Services provides clients with business process and industry expertise, a deep understanding of technology solutions that address specific industry issues, and the ability to design, build, and run those solutions in a way that delivers bottom-line value. To learn more visit: ibm.com

For more information:

Jonathan D. Breul

Executive Director

IBM Center for The Business of Government

600 14th Street NW

Second Floor

Washington, DC 20005

202-551-9342

website: www.businessofgovernment.org

e-mail: businessofgovernment@us.ibm.com

Stay connected with the
IBM Center on:



or, send us your name and
e-mail to receive our newsletters.