



Building Trust on the Internet

**An Intel and MIT study establishing
user trust in Web sites**

IT@Intel

Trust is one of the key factors in establishing positive relationships with customers on the Web. Intel and a team from the Center for eBusiness@MIT collaborated on experiments designed to explore how e-Business Web sites can become more effective and reduce support costs by increasing visitor trust in the site.

December 2003

Overview

- **The importance of Web trust**
- **Methodology—adaptive experimentation**
- **Experiments**
 - Trust seals
 - Decision assistance
 - Download wizard
 - Advisor persona
 - Advisor with voice
- **Benefits and key learnings**
- **Future work**
- **Conclusion**

IT@Intel

Intel IT and a team from the Center for eBusiness@MIT collaborated on a series of experiments designed to study the importance of Web trust. In this paper, we describe the project's background and discuss our methodology. We provide a synopsis of the experiments comprising the study, a summary of what we learned, and a description of what we are working on now and plan to work on next.

The Importance of Web Trust

- **Commercial success on the Internet depends upon the customer's specific experience on a site**
- **Can we reduce costs and build trust?**

IT@Intel

Trust is an important component of mutually beneficial customer/business relationships. Many types of e-Business Web sites—for example, customer support—now replace some or all of the interaction customers once handled by “live” personnel.

Successful e-Business Web sites reduce costs by allowing users to fill needs or resolve problems on the Web site rather than relying on a phone call. The success of such Web sites depends on customers' trust in the Web site.

Intel and MIT collaborated to study ways to establish trust on e-Business Web sites where customers do not have a direct interaction with a human respondent. We learned ways to reduce costs while building this trust.

Trust Affects Web Site Visitor Behavior

- **Desired behavior**
 - Download software
 - Buy a product or service
 - Enroll in marketing programs

IT@Intel

Our aim was to explore the factors affecting Web site behavior, especially how those factors affect Intel customers. More specifically, we wanted to explore the role trust plays at the point where a customer makes a decision, whether it be to download software, buy a product, or enroll in a market plan.

To accomplish this, we experimented on a full production Intel Web site dedicated to providing software downloads such as updates, drivers, plug-ins, and other support software.

Trust Factors

- **MIT has identified factors that contribute to trust**
 - Privacy and security
 - No coding errors
 - Brand
 - Touch and feel
 - Advice
 - Navigation
 - Fulfillment
 - Trust seals
 - and more...

IT@Intel

Our goal was to optimize our download Web site to strengthen trust. We looked at a variety of cues that help build trust, such as the look-and-feel of the site and the quality of the content (for example, the incidence of errors and the content's perceived value). We considered the impact of brand identity on trust, as well as the customer's perception of privacy and security.

We used nearly 100 factors identified by MIT as important to user trust in an e-Business relationship. These factors include privacy and security, error-free code, brand recognition, touch-and-feel, availability of advice, efficient navigation, successful fulfillment, and the display of trust seals.

Methodology

- **Adaptive experimentation**
 - Experiment
 - Update response model
 - Design new experiment
 - Update response model
 - Learn and continue improving
- **Key measurements**
 - Click-stream
 - Survey responses
- **Design**
 - Test and control
 - Before and after

IT@Intel

Our study was a collaboration between Intel and a team from MIT's Center for e-Business, led by Professor Glen L. Urban. Development and data analysis took place at Intel and at MIT. We conducted all experiments on the Intel public download site at downloadfinder.intel.com.

Using an adaptive experimentation methodology, we set up two environments: a control environment and a test environment. We assigned download site visitors randomly to either the control site or the test site, with about 80 percent sent to the control site and 20 percent to the test site.

We conducted a series of experiments, adding a new feature to the test site in each experiment. We investigated the impact of the additions by analyzing objective data obtained from click-streams, and subjective data gathered from user survey responses.

Throughout the sequence of experiments, we applied our learning to site improvement. We found that this helped us provide a highly useful interface because we could validate what was effective, and eliminate out what was not. Differences in success rates and overall satisfaction were determined by comparing test site data to control site data. We simultaneously obtained data from both the test and control

continued...

Methodology

- **Adaptive experimentation**
 - Experiment
 - Update response model
 - Design new experiment
 - Update response model
 - Learn and continue improving
- **Key measurements**
 - Click-stream
 - Survey responses
- **Design**
 - Test and control
 - Before and after

IT@Intel

sites time to control for external events that might alter user responses, for example a virus warning that might have made users reluctant to perform a download. In this way, we could determine whether the majority of users would benefit from an improvement before fully deploying it. We integrated the beneficial features into the control site, and repeated the cycle.

To obtain click-stream data, we started with a commercial click-stream logging tool that provided access to the raw log data, adding more capabilities as needed. This allowed us to analyze the user experience at a more granular level. We used a commercial survey tool to record user feedback from surveys.

We also collected data that allowed us to compare the effectiveness of the site before and after implementing a new feature. A key measurement we was “Did we answer your question?” This measurement used questions included in the user survey that asked, among other things, whether the user’s question had been answered or if the user had been able to successfully complete a task.

Experiment 1: Trust Seals

- **Objective**
 - **Validate the experimental presentation and data collection methods**
 - **Design test version control environments**
 - **Test impact of trust seals**
- **Overall results**
 - **Users noticed trust seals**
 - **Intel brand is effective**
 - **Slight increase in subjective security rating**
 - **Users perceived that the site addressed their needs (overall)**
 - **Added trust seals to live site after validation**

IT@Intel

The main goal of our first experiment, launched in early 2002, was to set up the infrastructure of the experiments. We created test and control environments that allowed us to anonymously identify repeat users, and established a quality assurance (QA) process to test every change made to each environment. We also incorporated click-stream analysis tools, and created a tailored a user survey to determine customer satisfaction.

In this experiment, users used a drop-down box to specify their operating system. Then they selected a file to download from a list of downloads available for the product they selected.

The difference between the test and control sites for this initial experiment was the presence of two trust seals. Although the focus of this experiment was not to measure the impact of trust logos, the user survey showed their presence brought a slight increase in user perception of the site's security, particularly among novice users.

Experiment 2: Download Assistance

- **Objective**

- Improve navigation
- Improve download completion and success rate
- Context-sensitive content

- **Overall results**

- 3% increase in successful downloads
- \$18 million reduction in annual customer service costs
- 18X return on investment in the first year
- Recovery over research costs
- 3.5% increase in trust (overall)
- 4.5% increase in customer satisfaction
- Decision assistance content added to live site



IT@Intel

In the second experiment we introduced context-sensitive download assistance to help the user make download decisions. We offered appropriate options depending on where the user was in the site and in the download process.

The objectives of this experiment were to improve the navigation flow through the site, improve completion rates (“Did users find what they wanted?”), increase the success rate of downloads, and implement this context-sensitive dynamic component throughout the site.

Based on user survey results, we saw a 3 percent increase in successful downloads. 3 percent may seem like a small improvement, but consider that approximately one million users each month use this download site. So a 3 percent increase means 30,000 additional downloads per month. A human response to an Intel customer costs around \$48 per call, so a 3 percent increase in successful downloads could potentially translate to \$18 million annual cost avoidance.

Other observations included a 3.5 percent increase in trust overall and a 2.3 percent increase in willingness to recommend the site. The most significant benefit was a 5 percent increase in

continued...

Experiment 2: Download Assistance

- **Objective**
 - Improve navigation
 - Improve download completion and success rate
 - Context-sensitive content
- **Overall results**
 - 3% increase in successful downloads
 - \$18 million reduction in annual customer service costs
 - 18X return on investment in the first year
 - Recovery over research costs
 - 3.5% increase in trust (overall)
 - 4.5% increase in customer satisfaction
 - Decision assistance content added to live site



customer satisfaction. We also found that users finished their task more quickly with download assistance than without assistance.

With the added decision assistance, the customers' satisfaction with getting their questions answered improved from less than 60 percent to over 80 percent.

Based on these results, we concluded that the context-sensitive download assistance component was effective and rolled it into the control site for the next release.

Experiment 3: Download Wizard

- **Objective**

- Provide alternative, step-by-step, user flow through wizard
- Tested only with cameras and graphics product categories
- Validate and compare user click-stream against surveyed responses

- **Overall results**

- 15% chose the camera wizard
- Wizard increased success rate 6.5% versus control
- Survey and click-stream data agree in direction, but not always in magnitude
- Download wizard added to control site
- Analysis is ongoing



IT@Intel

In the third experiment we introduced a download wizard that provided an alternative way for users to complete a download. Users could choose to use the decision assistance added in the previous step, or a download wizard that steps them through the download process.

To reduce implementation time and complexity, we implemented the wizard for just two of the product families: the PC cameras and graphics products. The user survey had shown that the users of these two product families were different: visitors to the PC camera product site were more likely to be novices; visitors to the graphics product site were generally more advanced. We wanted to see how each of these diverse populations would respond to the wizard.

We expected users who wanted a little more “hand holding” to select the download wizard option. However, results showed that less than 10 percent of visitors to the PC camera segment chose to use the wizard. Although this sample was small, it was still large enough to provide statistically significant data. Surprisingly, more visitors to the

continued...

Experiment 3: Download Wizard

- **Objective**
 - Provide alternative, step-by-step, user flow through wizard
 - Tested only with cameras and graphics product categories
 - Validate and compare user click-stream against surveyed responses
- **Overall results**
 - 15% chose the camera wizard
 - Wizard increased success rate 6.5% versus control
 - Survey and click-stream data agree in direction, but not always in magnitude
 - Download wizard added to control site
 - Analysis is ongoing



IT@Intel

graphics segment selected the wizard. This dispelled our expectation that more advanced users would be less likely to choose a wizard-type flow to assist them with downloading.

Use of the download wizard did increase download success. The analysis is ongoing, but so far indicates that the wizard has improved the site. Its apparent success led us to add the download wizard to the control site for the next experiment.

Experiment 4: Advisor Persona

- **Objective**

- Include human persona within existing wizard
- Step visitors through download process
- Test only with cameras and graphics product sections

- **Overall results**

- Higher success and satisfaction with advisor than without
- Advisor persona added to control site
- Analysis is ongoing



IT@Intel

In our fourth experiment we introduced our advisor persona, Rosa, to the download wizard. We added Rosa to both the PC camera and graphics segments so we could compare her effect on these two user groups.

With the advisor persona added, visitors to the graphics segment continued to choose the download wizard more than visitors to the PC camera segment.

We're still gathering and evaluating data, but the initial results indicate that satisfaction was higher with the advisor than without.

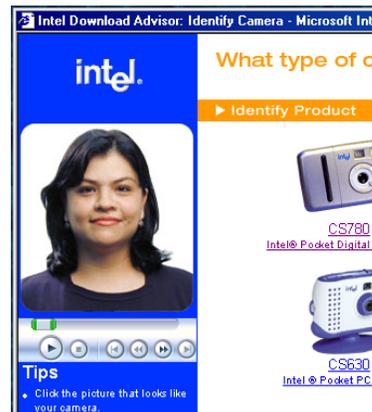
Experiment 5: Advisor with Voice

- **Objective**

- Introduce human voice to user experience
- Communicate user tips through voice output
- Dynamically offer expanded experience through client detection

- **Anticipated results**

- Experiment still underway
- Increase success rates
- Increase “Did we answer your questions?” satisfaction
- Enhance return on investment
- Verify that a more lifelike advisor enhances trust and is more effective



IT@Intel

In our final experiment we introduced audio to Rosa. Because the experiment site is a live production site with approximately one million users per month, it was a challenge to find an audio plug-in that worked for all users. We considered several commercial multimedia applications but in the end decided we didn't want to ask users to install a plug-in. We felt such a request might reduce the number of people willing to participate and affect experiment results. Instead, we found a way to embed an audio stream: using controls embedded in the operating system, we tapped the default multimedia application on the user's system.

We dynamically detected whether the user's browser security settings allowed download of the embedded controls, and if the browser could accommodate an audio stream. If not, we sent them to the control site, which used the version of Rosa without audio.

In this experiment we wanted to determine whether audio improved the “question answered satisfactorily” metric and the download success rate. We also wanted to determine whether a more lifelike advisor enhances trust and is more effective overall. Because we only recently launched this experiment, results are not complete.

Benefits and Key Learnings

- **Methodology: What we learned**
 - Establishing a test/control environment
 - Combining data from click-stream and surveys
 - Completing experiment phases with small deltas
 - Experiments run on live site with real users
- **Managerial**
 - Potential to improve even effective sites and improve self-help
 - Payback can be quite large – call deflection can add up quickly
 - Validation for high-end user interface elements and content offerings
 - Trust is an important strategic concept for collaboration and in strengthening customer relationships

IT@Intel

We found that conducting the experiments in phases, adding a single feature in each experiment, and using test and control environments allowed us to focus on features separately and better understand how each feature affects user trust and acceptance—before releasing it for general use. We learned that it was important to allow time and resources for quality assurance at every stage. We used rigorous QA for even relatively minor changes (such as adding the image of Rosa) before release in both the test or control environments.

Completing the experiment in small increments allowed us to more tightly control test conditions. We could validate the success of a new feature through click-stream and user survey data analysis before adding the change to the control site. When we did roll a feature out to a wider audience or a production site, we knew the feature had been validated with data supporting its effectiveness.

We learned that survey and click-stream data may not agree. A user might say in the survey that they were successful, but their click-stream showed otherwise. So we have developed a method to correlate subjective user survey responses with

continued...

Benefits and Key Learnings

- **Methodology: What we learned**
 - Establishing a test/control environment
 - Combining data from click-stream and surveys
 - Completing experiment phases with small deltas
 - Experiments run on live site with real users
- **Managerial**
 - Potential to improve even effective sites and improve self-help
 - Payback can be quite large – call deflection can add up quickly
 - Validation for high-end user interface elements and content offerings
 - Trust is an important strategic concept for collaboration and in strengthening customer relationships

IT@Intel

objective click-stream data. This helps us more clearly understand how the feature affects user trust and acceptance.

Running the experiments on a live site with real users gives us more confidence in our results. Live site experiments are effective tools for user analysis because we capture users' behavior in their natural environment—behind their desks, at work, at home. We don't invite users into a usability lab where they may not be performing tasks they would normally do or where they might feel uncomfortable or under pressure to perform.

Throughout the study, we asked people to join a long-term study. In this study we used unique, anonymous criteria to identify potential long-term participants. This lets us correlate the click-stream and survey results for an individual user each time they come to the site and correlate it to their behavior and survey responses during prior visits. We can then see how their experience changes over time.

Future

- **Complete analysis of experiments**
 - Tie click-stream analysis to survey responses
 - Longitudinal analysis of complex paths
 - Build model of trust formation to help in experimental design for opportunities
 - Rationalize differences between click-stream and survey
- **Potential extensions and future applications**
 - Wizard-like capabilities to assist customers in product identification on Intel's technical support Web sites
 - Address needs of Intel's distributors and partners
 - User-market investigation of product ID system prior to production launch
 - Spread learnings throughout www.intel.com
- **Trust and experimentation *can* payoff**

IT@Intel

Although we have analyzed the data enough to obtain significant results (such as the potential \$18 million in annual cost avoidance), we continue to analyze it to gain even deeper understanding. Some areas of focus include correlating the differences between click-stream results and survey results, analyzing usage patterns and level of satisfaction over time, and building a model of trust formation to serve as a basis for future experiments.

Our future work includes exploring and testing options for animating the advisor persona. Our test groups will be given incentives to encourage their continued participation in the long-term study that we hope will provide even greater insight.

Conclusion

- **Potential to improve even effective sites and improve self-help**
- **Payback can be quite large—call deflection can add up quickly**
- **Validation for high-end user interface elements and content offerings**
- **Trust is an important strategic concept for collaboration and in strengthening customer relationships**

IT@Intel

We determined that even effective Web sites can improve. Our site was considered effective, but with just a few minor changes we were able to increase user satisfaction, as evidenced by their greater than 80 percent satisfaction with their ability to get their questions answered.

The payback for these improvements can be quite large. At \$48 per service call, a 3 percent increase in successful downloads among 3 to 4 million users a month could result in an annual cost avoidance of \$18 million.

Our methodology proved to be an effective way to design high-end user interfaces. We learned that we could continually validate features and enhancements using nearly instant feedback obtained from click-streams and surveys.

This study has shown that trust is clearly an important strategic concept and one of the key factors in customer relationships.

Intel around the world

United States and Canada

Intel Corporation
Robert Noyce Bldg.
2200 Mission College Boulevard
P.O. Box 58119
Santa Clara, CA 95052-8119
USA
Phone:
General Information: (408) 765-8080
Customer Support: (800) 628-8686

Europe

Intel Corporation (UK) Ltd.
Pipers Way
Swindon
Wiltshire SN3 1RJ
UK
Phone:
England: +44 (0) 1793 403000
France: +33 (0) 1 58 87 71 71
Germany: +49 89 99143 0
Ireland: +353 1 606 7000
Israel: +972-4-8656060
Italy: +39 2 575 441
Netherlands: +31 (0) 20 6591800

Asia-Pacific

Intel Semiconductor Ltd.
32/F Two Pacific Place 88 Queensway,
Central
Hong Kong, SAR
Phone: +852 2844 4555

Japan

Intel Kabushiki Kaisha
P.O. Box 300-8603 Tsukuba-gakuen
5-6 Tokodai, Tsukuba-shi
300-2635 Ibaraki-ken
Japan
Phone: +81 298 47 8511

South America

Intel Semicondutores do Brazil
Avenida Dr. Chucri Zaidan, 940-10° andar
04583-904 Sao Paulo, SP
Brazil
Phone: +55 11 3365 5500



For more information

To learn more about IT@Intel, visit our site on the World Wide Web at www.intel.com/IT

This document and related materials and information are provided "as is" with no warranties, express or implied, including but not limited to any implied warranty of merchantability, fitness for a particular purpose, non-infringement of intellectual property rights, or any warranty otherwise arising out of any proposal, specification, or sample. Intel assumes no responsibility for any errors contained in this document and has no liabilities or obligations for any damages arising from or in connection with the use of this document.

Intel and the Intel logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

*Other names and brands may be claimed as the property of others.
Copyright © 2003 Intel Corporation. All rights reserved.