Peer-to-Peer Computing Enabled Collaboration

Martin G. Curley

Director, IT People, Intellectual Capital and Solutions,
Intel Corporation,
Leixlip, Co Kildare, Republic of Ireland
Martin.G.Curley@intel.com

Abstract. This paper discusses how peer-to-peer computing is emerging as a disruptive technology for global collaborative solutions. It explains how peer-to-peer computing can enable new collaborative solutions while significantly decreasing IT costs and improving IT asset utilization. An overview of the technology and usage models are discussed whilst the benefits are illustrated through a short case study from Intel. Finally the value proposition for peer-to-peer computing is summarized.

Introduction

Today's workforce faces unprecedented challenges and competitive pressures. Employees need to learn faster and colleagues often need to collaborate cross-functionally and globally, to create new products and achieve operational excellence.

In parallel, information technology is creating new value propositions, which can help businesses solve these problems. For example, Knowledge Management software is helping corporations to share and leverage knowledge better. The combination of motivated knowledge workers and applied information technology is creating what might be called a bionic organization—one that delivers turbocharged output, with the whole significantly greater than the sum of the parts.

Global Companies like Intel, which have sites distributed across the world rely on global computing and collaboration tools to conduct business. In fact it could be said that IT infrastructure and products and services are the air that modern day global enterprises breathe. There are many factors that influence the effectiveness of global corporations; key modulators include knowledge, skills, desire and collaborative friction. In fact one could postulate that an organization's potential output is somehow related to the product of these variables, i.e.

$$\text{Maximising Organization Output} = \text{Knowledge} * \text{Skills} * \text{Desire} * (1 - \text{Collaborative friction}) * \text{other relevant factors.} \quad (1)$$

To maximize output and quality from an organization, employees must have the right knowledge and skills to perform their jobs and they must have the ability to share and leverage knowledge easily. They of course must be highly motivated. Geographic and time zone separation are barriers which create collaborative friction. Leo Ellis captures a good example of low collaborative friction in his quote “*Engineers have yet to devise a better inter-office communication system than the water cooler*”. In many instances application of information technology is used in an attempt to reduce organizational collaborative friction. Peer-to-Peer computing is emerging as a very interesting computing platform to help lower collaborative friction.

Defining Peer-to-Peer Computing

In it's narrowest definition Peer-to-Peer Computing is fundamentally about the sharing of resources amongst computers. A broader definition of peer-to-peer is that it is an emerging class of applications, which leverages resources at the edge of a network. These resources include people, knowledge, information, network, computing power, storage. Whatever the definition, Peer-to-Peer computing is emerging as a disruptive technology for both collaboration applications and the underlying IT infrastructure.

It is a disruptive technology for collaboration applications as it is enabling the introduction of new capabilities which were previously cost-prohibitive to implement.

It is a disruptive technology for IT infrastructure as it can significantly increase the utilization of existing IT assets and improves the performance of supported applications.

Intel ® Share and Learn Software

Intel ® Share and Learn Software (INTEL SLS) is a good example of a peer-to-peer computing solution. It is essentially a peer-to-peer file sharing application, which helps solve multiple business problems while leveraging existing IT infrastructure.

INTEL SLS runs on corporate Intranets and using it's advanced content networking solution it can be used for applications such as eLearning and Knowledge Management. It delivers these capabilities while reducing information technology costs by moving network traffic off wide-area networks to much cheaper local area networks and by leveraging inexpensive PC storage. Using INTEL SLS each PC in an enterprise network can be converted to a caching device, significantly increasing performance for certain file transfer operations.

The foundational elements of INTEL SLS are a content index, and content publishing, management and distribution applications. A database server and web server provide a set of services to client users who use the web for retrieving and viewing content while they are connected to a corporate Network. They can use an offline player when they are not connected to the network.

INTEL SLS is content agnostic, meaning that it can manage and distribute a diverse type of files. It is particularly useful for the distribution of rich content files as it allows near seamless propagation of large files throughout a network using a peer-to-peer distribution capability.

The following sections discuss how INTEL SLS and peer-to-peer computing are enabling new collaboration capabilities at Intel Corporation.

Enabling New Collaboration Applications

Knowledge Management

One significant opportunity cost for corporations is that much codified knowledge exists in the form of presentations and text documents stored on individual's PC hard drives. In conventional networks this information and knowledge is not available to others for sharing and reuse.

INTEL SLS introduces a capability where any file can be shared with many others in a corporation, through a personal publishing function. The INTEL SLS Personal Publisher allows the filename and relevant metadata to be easily registered in a index which other users can search based on criteria such as subject, keywords etc. If a user wishes to retrieve the file, the system searches for the closest located copy of that file and initiates a peer transfer from the closest client. The file returned is always the latest version.

This is important, as in effect; an enterprise knowledge management system is being created, leveraging previously unavailable information and knowledge at the edge of a corporation. This creates new opportunities through knowledge sharing and re-use, whilst saving money through using in-expensive client storage. And uniquely to peer-to-peer solutions, the more the system is used the better the performance improves.

Intel itself is using INTEL SLS for its Intel Notebook of Knowledge (INK) program. This program captures codified IT knowledge in the form of white papers, videos and presentations. INTEL SLS using its peer-to-peer services acts as integrated object repository delivering a requested object from the closest client. Using INTEL SLS publisher, any employee can add content from their hard drive to the central index. All files are catalogued using the Dublin Core tagging standard.

e-Learning

The use of rich content, presented in a compelling and immersive fashion can significantly enhance learning effectiveness. Rich content uses technology to blend text, graphics, sound and motion to deliver

information. Enabled by new technology, rich content can deliver knowledge and information in a form that enhances understanding and retention by providing emphasis and context. This concept is well supported by the statement of Prof. Fred Hofsteter, Director, Instructional Technology Center, University of Delaware.

“People remember 20% of what they see, 30% of what they hear, 50% of what they see and hear and 80% of what they see, hear and do simultaneously”.

Couple rich content with the efficiency of e-Learning, where content is recorded once and used many times (instead of the each-time requirement for facilities, trainer and materials of conventional training) and it creates a win-win scenario. INTEL SLS enables the delivery of rich content and large compressed video files almost seamlessly over an enterprise network.

The Intel IT organization recently converted its new hire integration program to a fully virtual e-Learning program, delivered to new Intel IT employees using Intel SLS. The integration program consists of sixteen modules including technical, business, customer service and interpersonal skills, which new employees take over a six-week period. The net result was that integration satisfaction survey scores improved to 88% from 77% whilst significant travel and training cost savings were achieved. The content is reusable and other Intel IT employees are taking these classes at no extra cost.

Corporate Communications

The use of Oral History is becoming increasingly important in Corporate Communications. Recording a short video clip is often a more effective and efficient method of communication rather than taking the time to write down content and then distribute it via mail or paper memo. Dave Snowden, Director of the Institute of Knowledge Management in EMEA says *“I always know more than I can tell, and I always know more than I can write down”*. Hence large corporations are increasingly using short video clips for transfer of key knowledge/messages to employees worldwide.

INTEL SLS introduces the capability to distribute these video clips almost seamlessly throughout the enterprise with no additional bandwidth or infrastructure required. 20-40 MB files download in less than a minute rather than the 20 minutes, which you might typically see across a wide area network. This performance is achieved through the local caching that INTEL SLS provides, thus avoiding multiple wide area network downloads.

Intel Senior Vice-presidents like Mike Splinter and Sean Maloney are using INTEL SLS as another channel to deliver monthly video updates to their teams worldwide. The video quality is close to broadcast quality and conveys key monthly messages and challenges in short clips.

The Corporate IT infrastructure equation

At a macro level IT computing infrastructure costs in a large corporation can be described as the sum of the costs of providing the compute-network-storage platform as well as the costs of managing this environment.

Corporate IT Platform (Performance & Cost) = Sum (Computing, Storage, Network, Management) Performance & Costs (2)

A primary function of a Corporate IT organization is to optimize this equation to ensure the best performing corporate computing platform is delivered at the lowest possible cost. For Corporate IT organizations the cost of providing and supporting the enterprise platform continues to escalate which users continually demand better solutions and performance. These user and Corporate IT goals are often conflicting.

Using peer-to-peer computing protocols, utilization of existing corporate IT assets can be increased and equivalent or better performance can be delivered at lower cost. This also allows implementing of new capabilities faster and cheaper than implementing new systems. This is essentially due to the fact that Peer-to-Peer computing allows the substitution of cheaper components for more expensive computing components in the Corporate Infrastructure equation. In-expensive client hard-drives can be used for storing

information instead of more expensive network or server attached storage, Computing power on high performance PC's can be used instead of more expensive Server computing power and network traffic can be shifted from expensive Wide Area Networks (WAN) to Local Area Networks (LAN). Peer-to-Peer Computing will likely co-exist as a complimentary computing platform to client-server.

In order to assess the possible impacts of Peer-to-Peer Computing on corporate infrastructure, Intel's IT organization ran a number of internal trials within Intel.

IT trials results

In conjunction with the Intel training organization, Intel ran an internal trial using an eLearning class called American Accent, which consisted of about 60 modules of file sizes ranging from 5 MB to 15 Mbytes. In a trial with 1900 participants in over 50 distributed sites, users of this eLearning saw an average improvement of 4-5X in file delivery times whilst over 80% of the 3000+ files transferred were peer-to-peer transfers which took place over the LAN rather than the WAN. Figure one shows a comparison of transfer times for a large file using a peer-to-peer application compared to a traditional client-server environment.

A complimentary driver for the adoption of peer-to-peer computing is the relentless march of Moore's law. Instead of streaming video over a network, an alternative is to compress a video, perform a peer-to-peer transfer, decompress the video and then play it locally on a PC. As computing power increases and CPU architectures become smarter, decompression times are decreasing rapidly – for example a compressed 40 Mbyte video file that took approx 40 seconds to decompress on a 700 MHZ Pentium ® III processor based PC now takes approx 5 seconds to decompress on a new 2.2 GHz Pentium ® 4 processor based PC.

In an another experiment with the distribution of a video speech by Intel's CEO to employees worldwide, multiple thousands of copies of this video were distributed to employees worldwide with greater than 90% of the transfers being peer-to-peer, significantly decreasing file delivery time while lowering networking costs.

Based on the trial results Intel is aggressively adopting Peer-to-Peer computing solutions for collaboration solutions.

Peer-to-Peer Value Proposition

Peer to Peer computing is emerging as a great opportunity for enterprises to take advantage of latent potential in enterprise infrastructure whilst delivering new capabilities and better performance. The adoption of Peer-to-Peer computing drives benefits in the following area.

- Improved Performance
- Lower IT Infrastructure Cost
- Improved IT Asset Utilization
- Smarter Dynamics of Enterprise Systems
- Improved Fault tolerance and scalability
- Ability for more flexible deployment

As the security models of peer-to-peer computing solutions improve, this new computing model is destined to play a significant role in global collaborative solutions in the Twenty First Century.

Conclusion

The new computing model that the peer-to-peer era is ushering in, creates opportunities to enhance collaborative solutions and thus lower collaborative friction to help deliver improved organization performance. In addition the ability of the peer-to-peer computing model to allow better integration of different computing components to deliver improved performance, lower platform costs and better IT asset utilization make it rather unique in terms of computing models. In fact, peer-to-peer computing just may be one of the best-kept computing secrets of the late Twentieth Century.