Networking Metaphors for E-Commerce

Erik Wilde[1] and Andreas Steiner[2]

[1]ETH Zürich (Swiss Federal Institute of Technology)
[2]eTrade Solutions, Swisscom IT Services AG

TIK Report 190 (February 2004)

Available at http://dret.net/netdret/publications#wil04f

Abstract

E-commerce technologies have reached a level of maturity where many businesses are no longer hampered by technological limitations. However, the adoption of e-commerce technologies is slower than anticipated. We argue one of the limitations is a psychological barrier, which is created by the perception that e-commerce technologies are a whole new set of technologies which are completely different from computer networking. By applying metaphors from basic networking technologies (such as bridges and routers), we try to (1) demonstrate that e-commerce technologies are — in many ways — comparable to computer networking, and (2) show that convincing businesses to adopt e-commerce technologies could be made easier by showing them that e-commerce is basically computer networking taken to another level. We also believe that using these metaphors will make it easier to talk about e-commerce technologies, to reuse existing knowledge about networking architectures on this new level, and to identify the areas where additional work needs to be done.

1 Introduction

Basic e-commerce technologies have reached a level of sophistication which is sufficient to design and implement the majority of today's business processes. However, several studies (for example by *Berlecon Research*[1] or *Intellect*[2].) have shown that e-commerce adoption is slower that anticipated. Even though the purely technical barriers have become rather low, in many cases other factors seems to slow down the development which often is characterized as leading from a traditional model via *e-commerce* to the fully transformed

[1] Available online at http://www.berlecon.de/output/studien.php?we_objectID=125.
[2] Available online at http://www.intellectuk.org/publications/reports.

e-business. For our purposes, the distinction between the latter two is unnecessary and we will use the term e-commerce to refer to all IT-based means for doing business.

One of the more important reasons why e-commerce adoption is slower than anticipated is the *weak economy*. The economic developments of the past years has severely reduced the willingness (and ability) to invest. Investing in technologies that touch the very heart of a company has been delayed except in cases where change was inevitable, for example in the case of suppliers of major industries. In these cases, the buyers (being large companies) demanded fully automated supply chains for better integration into their business processes. The suppliers had to implement these supply chains in order not to lose their contracts. In many other cases, however, streamlining and redesigning the business processes has been more of a long-term ideal than an unescapable short-term goal. The recent economic upturns will probably improve the willingness to invest in e-commerce, but at the time of writing, this is more a vision than actual observation.

Another issue which hinders e-commerce adoption is what could be called the *dot.com backlash*. The so-called "new economy" caused a large number of catastrophic failures among the companies investing a lot in advanced technology. Strictly speaking, this did not prove that something is fundamentally wrong with new technologies, but it created an atmosphere where more conservative approaches to doing business have regained some appeal. Furthermore, the promises communicated by dot.com advocates sometimes were much too utopian, with one of the common misconceptions being that e-commerce and Web Services will enable buyers and sellers to search for each other on the Internet and cooperate through fully automated mechanisms. This vision will not become reality for a long time to come (and maybe it never will), and the discovery that some of the early promises went much too far has also contributed to the hesitation of companies in e-commerce adoption.

The single most important technical hurdle is *security*. Many companies feel that security is still a a problem. Exposing access to critical business processes certainly has a high potential for misuse, and many companies are uneasy about defining and publishing interfaces to their business processes. Security issues are increasingly addressed by technological means (in particular the work of the W3C's *XML Encryption WG*[3]). However, we also believe that the security issue is typical for something that is also psychological, the unwillingness to adopt something that is not perceived as being fully understood, in general use, and generally accepted as safe.

This is what we call *high-tech angst*, something that for a long time kept companies from connecting to the Internet. The idea of having data packets from all over the world being sent into the corporate network was very troubling for IT professionals who were used to closed LANs. The development of networking protocols and devices such as HTTPS and firewalls has made it much easier to deal with the Internet, not only from the technological side, but also from the psychological point of view. Today, only very rarely people would argue to not connect a company to the Internet because of security issues. It is common wisdom that a properly configured firewall in most cases will be sufficient to provide a connection to the Internet where the benefits outweigh the risks.

[3] The group's home page is `http://www.w3.org/Encryption/2001/`.

Our assumption is that e-commerce technologies still struggle to reach this level of being conceived as a natural thing to do, with a number of standard mechanisms that are tested and trusted by the majority of potential users. Ten years ago, the Internet infrastructure that today is ubiquitous and taken for granted, had to overcome the same obstacles. In this paper, we propose to revisit the development of networking, and reuse the know-how and confidence that we have gained over the last ten years. In Section 2 we describe some of the basic networking concepts in use today, and some of the application-oriented networking devices that are the backbone of modern computer networks. In Section 3 we then apply these concepts and devices to e-commerce by interpreting them as metaphors for e-commerce infrastructures.

2 Networking Metaphors

In this section, we describe the metaphors that we are using in our comparison of e-commerce infrastructures and traditional computer networking. We then continue in Section 3 by applying these metaphors to e-commerce infrastructures and investigate the overall picture as well as some case studies.

Computer networks are a commodity nowadays, and the structure of computer networks is well-known and has been studied for a long time. Today's Internet-oriented network architectures distinguish four basic kinds of devices inside a network, where the classification is based on the protocol layer:

- *Repeaters:* Repeaters operate on layer 1, the physical layer. Basically, repeaters amplify electrical (or, for fibre-based networks, optical) signals, for example for the purpose of using longer network cables than normally possible (because of a cable's inherent attenuation). Repeaters do not attempt to interpret a signal in any way, their only purpose is to read it, and then to amplify (or maybe regenerate) the exact same signal for further distribution. Another popular name for a repeater is *hub*.

- *Bridges:* A bridge is a layer 2 (data link layer) device. A bridge is smart enough to selectively decide which packets to forward to a particular network segment (often analogous to some interface of the bridge). A good example for a bridge is a *Wireless LAN (WLAN)* access point[4], forwarding packets from a LAN cable to the wireless interface and vice versa. The bridge will only forward those packets it receives on its wireless interface that should be forwarded through its cable interface. Another popular name for a bridge is *switch*.

- *Routers:* While bridges operate with MAC addresses (which identify the physical interface of a networking node), routers operate on layer 3, the network layer. In case of the Internet, this means that routers operate on IP addresses. In order to

[4] WLAN access points can also be configured as routers, implementing additional services that are not on the bridge level, such as *Network Address Translation (NAT)* or the *Dynamic Host Configuration Protocol (DHCP)*.

forward an IP datagram to the right interface (and finally to the right end system), routers need to receive information via routing protocols, which enable them to make the decisions that are required to make global networking possible.

- *Gateways:* Anything that operates above layer 3, but still is considered to be part of the networking infrastructure, is a gateway. Gateways can have very different tasks, they can transfer data between different networking architectures, or they can implement application-specific functionality that requires knowledge of particular application-level protocols. A typical example for a gateway is a *Wireless Application Protocol (WAP)* gateway, that connects the Internet and a mobile phone network (such as a GSM network) and enables mobile phone users to retrieve Web pages with their mobile devices.

The main lesson from these networking devices is that a classification of devices is useful and helps to build networks out of a given set of networking elements. In the networking world, the network model itself, using the layered structure, has become the globally accepted way to model computer networks. This structured view of networks makes it possible to distinguish between devices that can be used for different tasks inside a network.

In the late 80's, there was an ongoing debate whether the global computer network should be built on top of the Internet protocols or on top of the *Open Systems Interconnections (OSI)* stack of protocols [8]. The OSI stack was much more complex, and finally implementation difficulties with some OSI protocols and the overwhelming complexity of the OSI stack led to the adoption of the Internet protocols. From the architecture point of view, the basic network layers were the same (layers 1 through 4), but the OSI stack also defined two more application oriented layers:

- *Session Layer (Layer 5):* The session layer provides the control structure for managing end-to-end communications. Possible examples are establishing, managing, and terminating sessions. The rationale for the session layer is that many applications require communications beyond the simple and stateless transfer of data packets.

  In the Internet world, session management is managed by applications. In the case of the popular *Hypertext Transfer Protocol (HTTP)*, the underlying Internet transport protocol does not provide session management. Consequently, session management must be layered on top of the application protocol (popular techniques for HTTP session management are Cookies and URI rewriting).

- *Presentation Layer (Layer 6):* The transport layer provides platform-independent ways of encoding data. It includes mechanisms for communicating peers to negotiate a transfer encoding, and support to produce and consume this transfer encoding. The most important standard on this layer is the *Abstract Syntax Notation One (ASN.1)*, which is still in use today in various application areas (X.509 certificates are a popular example).

> In the Internet world, XML has taken on the role of a globally accepted transfer syntax. Even though there is no negotiation mechanism (which is not necessary, because there is only one syntax for XML) and no established view of the underlying information model [10], XML certainly has filled a gap that has hampered many interesting developments in global networking of applications.

Since the OSI layers 5 and 6 have no counterparts in the Internet protocol world, there are no established metaphors for networking nodes working on these layers. However, there have been developments in these areas, even though they have not been widespread enough to become a generally accepted networking node. An example for layer 5 support are Web application development frameworks providing support for session management. These frameworks provide programmers with the notion of sessions, and programmers are free to build their applications as if Client/Server-relationship in the Web were stateful. The frameworks then can be configured to implement the session support by using either Cookies or URI rewriting, and they can make this decision even dynamically depending on whether a particular client provides support for Cookies or not.

An example for a typical layer 6 application is a WAP gateway transcoding text-based *Wireless Markup Language (WML)* pages into *Wireless Binary XML (WBXML)*, which is an more compact form of WML which is better suited for the low bandwidth of today's mobile networks. In this case, the server sends the WML page in its text-based syntax, but the client (the mobile device) receives the page in a more compact binary syntax. The idea behind this model is that the content of WML is not altered in any way, it is only the encoding of the page that changes.

Apart from these rather specific examples, there are at least two examples of networking nodes that have become ubiquitous in today's networking world. Nodes of these types can be found in almost every network configuration, but the actual function of these nodes heavily depends on their configuration.

- *Proxies:* A proxy is a network node that accepts requests on behalf of another server, and either replies to the request (for example, because the local cache could be used), or forwards the request to another server (which can also be a proxy, in which case there is a chain of proxies). Proxies not necessarily contain caches, but many proxies are set up to include caching functionality.

- *Firewalls:* A firewall is a network node that selectively forwards or rejects packets and/or requests. Firewalls can operate on different layers of the networking protocol stack. On the lower layers, firewalls make their decisions based on Ethernet MAC addresses or IP addresses, for example rejecting all packets on a LAN that do not originate from a MAC address that is known to the firewall. Higher layer firewalls may work on application protocols such as the *Simple Mail Transfer Protocol (SMTP)* or HTTP. They may even attempt to interpret the messages being sent through application level protocols, so that essentially a *spam filter*, rejecting unsolicited e-mail messages, also can be regarded as a firewall.

The metaphors introduced so far cover the vast majority of networking node types being in use today. While there certainly are some exceptions, most nodes can be classified into one of these categories. This makes it easier for network architects to design, manage, and extend a network. Because e-commerce works on a higher level, it normally requires more sophisticated processing than raw network data (where most of the processing is done based on packet or datagram formats that allow only little variation of the values), and the following development offers an interesting analogy in this area.

As a more recent development in networking, in the early 90's the topic of *Active Networks* [9, 2] emerged among computer communications researchers. Active networks are networks where the network nodes do more than simple routing. The idea of an active network is that the network nodes can be configured dynamically to perform application-specific tasks. One possible example is video streaming, where an active network node (a router) should be able to downscale a video stream by transcoding it, because the router may know that the downstream connection is too narrow to carry the full video signal. Thus. active network systems must have some way of installing executable code on network nodes, which obviously poses problems of security and efficiency.

In principle, proxies and firewalls could be regarded as special cases of active network nodes, in the sense that a fully programmable router could be programmed to act as a proxy or firewall. However, many active network architectures have boundaries for the execution time and function libraries (such as secondary memory access, which is essential for the caching that a typical proxy performs), so that dedicated network nodes with a well-defined functionality will remain in use for some time to come. Active networks are more a research topic than a deployed networking technology, but the concept itself is very promising and the following section shows that the architecture of an active network can be of use on the e-commerce level.

3 E-Commerce Applications

E-commerce as it is viewed today most of the time builds on top of *Web Services*, which provide support for distributed programming in a heterogeneous environment. As the basic means of transport, the *Simple Object Access Protocol (SOAP)* [1] is used to transport web service requests and (optionally) responses. SOAP is a message format independent from a particular transport mechanism, but its design is based on an end-to-end communication model. Viewed from this perspective, e-commerce looks like a typical end-to-end architecture. However, in many scenarios today, SOAP is not used as an end-to-end protocol between business partners (like the buyers and sellers of goods), but as a protocol between one business partner and a central entity, the *electronic marketplace*. Rather than looking at electronic marketplaces as applications which are full-blown communication peers in the e-commerce architecture, i could make more sense to look at them as intermediaries, performing format conversion and possibly some other services. In Section 3.1 we take a closer look at these intermediaries and why they could be regarded as gateways when referring to networking metaphors.

While the gateway analogy is useful for thinking about more efficient and flexible formation of e-commerce networks, it does not address security issues. In our examples, we look at security from two different points of view. The first perspective deals with security in the sense of data integrity and quality. In many business process flows, data quality is constantly decreasing, because every processing step may introduce errors and typically does not reestablish data quality loss that may have been caused by previous processing steps. As one possible solution to this problem, data quality firewalls are discussed in Section 3.2, which have the same purpose as network firewalls, but use different technologies to perform the filtering functions usually associated with firewalls.

The second perspective on security is more target towards protecting data. Section 3.3 discusses what in a normal network would be a gateway providing authentication and/or encryption to another gateway, effectively establishing a secure tunnel to a communications peer. This example concludes our list of more detailed case studies of e-commerce network nodes. In Section 4, we discuss some more possibilities to further exploit the analogy of networking metaphors and e-commerce networks.

3.1 Conversion Gateways

One of the main challenges in e-procurement is the conversion of different file formats, such as files in any of the numerous e-commerce related XML standards, into another, or converting EDIFACT messages into XML (or vice versa). Converting an XML file adhering to some standard into another XML-based standard often can be achieved by using XSLT code. However, there is more to file conversion than simply map one format to another. Usually, the different standards use different content, or — in other words — a database of one trader in a trading community usually contains 80% standard information and 20% company specific information. If one trading partner wants to exchange information electronically with another one, the information exchange needs to be in sync, meaning that both have to agree on very detailed level what data is exchanged (apart from the format or container they use).

Nowadays, electronic marketplaces close this gap by allowing a trading partner to connect once to the marketplace using any file format and interface suited to the infrastructure of the trading partner. The marketplace has to verify that the data exchanged between trading partners is in sync, and it transforms the data from one file format to the other. Additional services like data enrichment etc. can be provided by the marketplace.

Basically, marketplaces add value in a trading network based on electronic data exchange by transforming orders, invoices etc from one file format into another, and dealing with aspects like sparsity (one trading partner cannot provide the data the other trading partner expects) and data overflow (data needed by one trading partner in the whole e-procurement process, but the other trading cannot deal with some of the data because his infrastructure cannot store and retrieve the additional data).

Thus, marketplaces can been seen as conversion gateways meaning that they deal with all the aspects involved in exchanging information from one file format with a specific data content into another format (with slightly different content). The need for such conversion gateways is due to the variety of standards for data exchange and additional non-standard solutions, as well as the need of companies to exchange company specific data during the procurement process and have access to it after completion, such as company specific account information in an invoice.

Since there will never be a single standard for documents in an e-procurement process, the format conversion and all other issues regarding the conversion will remain. To resolve the issues and being able to provide a flexible and dynamic solution, the idea of a conversion gateway can be seen as a part of the still evolving e-commerce infrastructure. If a node in such a network is capable not only of routing messages, but also of converting documents from one format to another (depending on the sender and the receiver of the document), then any service provider in the network could add additional value to the node. Additionally, this would probably help to limit the amount of standards supported and would help to focus on the business issues rather than reinventing the wheel again and again by constantly coming up with new software architectures for marketplaces.

3.2 Data Quality Firewalls

Data quality has many facets. In the e-commerce world governed by XML, on the most basic level it means that data should correspond to some schema, with the most popular schema languages being *Document Type Definition (DTD)* and *XML Schema*. While DTDs only guarantee the structural integrity of a document, XML Schema also supports a sophisticated datatype concept. In both cases, the correctness of a given XML document is tested for by validating the document against the schema. A more flexible approach to schema validation is the *Document Schema Definition Languages (DSDL)* framework, which is based on a modular way to combine different schemas for testing for different facets of data quality. Other solutions have been proposed, such as the system described by Nentwich et al. [5]. Some of the more sophisticated solutions for validation provide support for accessing external data, so that an XML document can be validated against the content in a database system.

On a more abstract level, Riggs [7] proposes the concept of data quality firewalls, which are strategically placed nodes in an e-commerce network. These nodes can either be placed on a company's boundaries, to ensure that only valid e-commerce messages will be accepted, or on organizational boundaries within a company, for example between different departments. This way, the business process steps within a company would be required to maintain a certain level of data quality, because otherwise the messages would not be accepted by the firewall. Naturally, the filter functions employed by data quality firewalls would be different from the filter functions performed by network level firewalls. The functions would range from standard XML validation (using DTD, XML Schema, or DSDL) to very company-specific logic, which probably would be hard to implement declaratively.

However, the important aspect is to view data quality firewalls as dedicated devices which perform filter functions only (i.e., they do not change messages in any way). They should be configured using as much declarative code as possible (using schema languages), and probably some amount on non-declarative code, which ideally should be written in a platform-independent language, such as Java or XSLT. While e-commerce firewalls would probably in many cases require some executable code, substantial amounts of code should be declarative, making it much easier and more cost-effective to author, maintain, and extend.

3.3 Security Domains

A typical configuration of today's computer networks inside companies is a so-called *intranet*, which is a self-contained connected to the global Internet through some kind of access technology. In most cases, the intranet is shielded from the Internet through a firewall, which is configured to block unwanted traffic. Very often, the intranet is treated as a trusted network, usually referred to as a *Demilitarized Zone (DMZ)*, the assumption being that access to the intranet is limited to corporate computers only, and these are known not to be potential attackers. Even though this setup is not always sufficient (because of security holes in computers as well as networks), from an economic point of view it makes sense to concentrate security mechanisms at the gateway where data exchange between the intranet and the Internet is concentrating.

In the Web Services world, most architectures are build on top of end-to-end assumptions, and XML-Encryption technologies provide means to encrypt and/or sign Web Service communications. An interesting proposal for securing Web Service communications in a more flexible way has been proposed by Melzer and Jeckle [4]. The mechanism presented by them is similar to the network technology of a *Virtual Private Network (VPN)*, which also uses encryption mechanisms such as the *IP Security Protocol (IPsec)* or the *Point to Point Tunneling Protocol (PPTP)* to enable secure communications over an untrusted network. VPN gateways are used to connect isolated parts of intranets, they facilitate the view of a single intranet, even though the underlying network infrastructure contains insecure and untrusted network segments.

Compared to the more usual end-to-end approach, a gateway-based Web Service security architecture has the advantage that encryption only takes place if necessary. If Web Service messages are using the intranet only, they remain unencrypted, which avoids processing overhead. More importantly, the gateway-based approach makes it possible to concentrate security-related tasks to the gateways, which results in better control over important data such as certificates, and makes deployment of security-related data such as updated certificates or certificate revocation lists much easier.

4 Further Work

In the preceding section, some case studies have shown that it is possible to identify e-commerce network nodes, and to establish a new way of thinking about e-commerce infrastructures. The majority of users today views e-commerce infrastructures that is complex, costly to develop, expensive to deploy, and hard to maintain. Contrary to this view, computer network infrastructures have reached the status of commodities, which consist of off-the-shelf components that are easy to configure and maintain. The long-term goal of e-commerce research has to have this picture of network components as the primary goal in mind. In this paper, we have not given any examples for e-commerce repeaters and bridges, but taking the analogies of computer networks, the following possibilities spring to mind:

1. *Repeaters:* These are the least intelligent devices in network architectures. However, in an e-commerce architecture based on a broadcast model, e-commerce repeaters could provide useful services. A possible scenario is a company implementing a model of Web Services where the services are distributed and may be provided at any place within the company's network. E-commerce repeaters could then broadcast Web Service requests inside the company, and the any service provider able and willing to provide the service could then reply to the request.

2. *Bridges:* Bridges could provide services similar to the repeaters discussed in the previous paragraph, but they could be placed at strategic places within an e-commerce infrastructure, for example boundaries between a company's departments. Inside the departments, a broadcast-based model based on repeaters could be used, while inter-department messages would be forwarded by e-commerce bridges.

While more end-to-end oriented e-commerce infrastructures as promoted today probably do not require e-commerce repeaters or bridges, it is possible that alternative models of using Web Services will create a demand for this kind of device. The architecture described by Chesholm et al. [3] is one example for a way of using Web Services which could expand today's prevalent way of looking at e-commerce infrastructures.

5 Conclusions

E-commerce is still in its early stages, and while the adoption rate is promising, it is lower than originally expected. One of the reasons is that potential users view e-commerce as a new and complex set of technologies, and very often this attitude is combined with security concerns. By establishing well-known computer networking metaphors for e-commerce infrastructures, we can lower the barrier to entry on the psychological level. This way we can also communicate that e-commerce infrastructures simply are a more application-oriented kind of computer networking that can benefit from existing know-how and technologies. We also think that we make exploring new opportunities easier by exploiting the common characteristics of computer and e-commerce networks, and concentrating on the parts where e-commerce networking is different and needs new approaches and solutions.

On a more technological level, research areas such as active networks, which very often are regarded as technologies which are only relevant for computer networking on the lower levels, could be used to learn from existing results and systems. While some of the use cases and constraints certainly are different for computer network level active networks and the application of similar technologies for e-commerce infrastructures, there also is a lot of overlap, and a deeper understanding of these areas would make it easier to improve e-commerce infrastructures by reusing existing knowledge and experience from networking infrastructure design.

References

[1] Don Box, David Ehnebuske, Gopal Kakivaya, Andrew Layman, Noah Mendelsohn, Henrik Frystyk Nielsen, Satish Thatte, and Dave Winer. Simple Object Access Protocol (SOAP) 1.1. World Wide Web Consortium, Note NOTE-SOAP-20000508, May 2000.

[2] Andrew T. Campbell, Herman G. De Meer, Michael E. Kounavis, Kazuho Miki, John B. Vicente, and Daniel Villela. A Survey of Programmable Networks. *ACM Computer Communications Review*, 29(2):7–23, April 1999.

[3] John Chelsom, Stephen Katz, Andrea Zisman, and Ron Summers. Information Bus. In Proceedings of XML Europe 2003 [6].

[4] Ingo Melzer and Mario Jeckle. A Signing Proxy for Web Services Security. In Robert Tolksdorf and Rainer Eckstein, editors, *Berliner XML Tage 2003*, pages 292–304, Berlin, Germany, October 2003.

[5] Christian Nentwich, Wolfgang Emmerich, Anthony Finkelstein, and Ernst Ellmer. Flexible Consistency Checking. *ACM Transactions on Software Engineering and Methodology*, 12(1):28–63, January 2003.

[6] *Proceedings of XML Europe 2003*, London, UK, May 2003.

[7] Simon Riggs. Data Quality and XML Validation. In Proceedings of XML Europe 2003 [6].

[8] Marshall T. Rose. *The Open Book*. Prentice-Hall, Englewood Cliffs, New Jersey, 1990.

[9] David L. Tennenhouse, Jonathan M. Smith, W. David Sincoskie, David J. Wetherall, and Gary J. Minden. A Survey of Active Network Research. *IEEE Communications Magazine*, 35(1):80–86, 1997.

[10] Erik Wilde. XML Technologies Dissected. *IEEE Internet Computing*, 7(5):74–78, September 2003.