Second Edition *including the story of USB 3.0, and the introduction to the author's next book:* un-kill creativity

bowling with a crystal ball

how to predict technology trends, create disruptive implementations and navigate them through industry

Yoram Solomon, PhD

18. Putting it all together: The story of USB 3.0

Timing is everything… This chapter was originally included in the manuscript of this book. However, as the manuscript was submitted for editing in May of 2007 it was not clear yet when will Intel announce the formation of the USB 3.0 promoters group, the first step towards releasing a new USB specification. As a result, and not to violate confidentiality, I removed the chapter and released this book in July. Sure enough, Intel announced the formation of the Promoters group two months later, in September. This chapter was important enough to release a second edition of the book, as it brings everything together: from technology trend analysis, to identifying a disruption opportunity (albeit not a major one) and its value proposition, to navigating it through industry politics. The USB 3.0 story is the story of all of those, and is the motivation for writing this book, even though it wasn't included in its first edition.

My writing style in this chapter might be slightly different than in the other chapters of this book for two reasons. First, because this chapter was written eight years after the other chapters. Second, because this chapter is more historical, and tells a story, rather than describes a methodology and theory. I chose to write it in a style somewhat similar to that of a Harvard Business School Case, except that it is told in first person. But enough conjecture—it is time to tell the story.

An Individual Contributor

In August of 2004 I was leading the *Consumer Electronics Connectivity* business unit at Texas Instrument. It was close to $100 million business unit, responsible for all silicon components concerned with connecting consumer electronics devices to one another. Those components used DVI, HDMI, FireWire, and other connectivity standards. And USB. Of course. Except that we (Texas Instruments) were infamous for being always behind with anything related to USB. There was a well-known "USB aversion" at TI. Some of the strongest FireWire advocates (and USB opponents) were in my group. We built components that connected iPods with PCs, digital camcorders with TVs, and the like. There were 89 people in my group, including myself. I can't honestly say that I enjoyed the operational side of being a business unit general manager, and in August 2004 my supervisor sent me to attend a week-long training at the *Center for Creative Leadership* in Colorado Springs. That was an amazing week that left a mark on me, one that will eventually lead to the creation of *Large Scale Creativity* in 2015. Throughout the week we examined and dissected my leadership skills and capabilities. I interacted with leaders from the top companies in the world. On the last day we each met with an organizational psychologist. As I sat down with her, after she had reviewed my resume and professional history, she wanted to frame our conversation around: "are you a *startup* person or a large company person?" Given my experience with startups, she was surprised to learn that I now worked for a Fortune 500 company with more than 35,000 employees. But there was another question that was burning in my mind, and she agreed to frame our discussion around: "am I a leader, or an individual contributor?" We talked for the best part of two hours, after which I attended the graduation ceremony, and boarded the plane back home. It is a relatively short flight from Colorado Springs to Dallas, but when we landed in Dallas, I knew exactly what I had to do Monday morning. The next week I asked my supervisor for a personal discussion. I told her that I felt more as an individual contributor, a strategist, rather than as a team leader. I asked if I could move from leading the Consumer Electronics Connectivity business unit to become the strategist for the group of three business units (one of which was the CEC). This conversation didn't end there, and we continued to discuss that over the next several months.

That September, like every year, we started the process of setting the priorities for the following year, 2005. The connectivity standards that our products supported were maturing and commoditizing, and as a result our gross profit margins were declining. Revenue was important, and all three general managers were asked to include a "top-down" priority of *increasing revenue 20% over 2014*. Something was missing, I felt, and took it me a couple of days to put my finger on it. I met with my supervisor again and told her: "I can see that we want to increase revenue 20% year-over-year, but how about a goal of increasing revenue 100% in 4 years?" To that she replied: "it's the same things. If you grow revenue 20% every year, then in 4 years you would have grown revenue by 100%." Actually, 107%… But to me those were two completely different things. In order to grow the unit's revenue by 20% in one year, the products you have to sell should already be in existence. This could only be done by selling more of the same products to the same customers, or sell the same products to new customers. It could not be done through the development of new products, because this would take much more than a year. Especially in the semiconductor industry, even minor changes to a product would take more than a year to release. If we only focus on a one-year horizon, we will never develop new products, and with the maturity and commoditization of our existing products—growing revenue will get harder and harder. “I disagree” I replied to her, “growing revenue 20% a year requires turning and pulling a different set of knobs and levers than growing revenue 100% in 4 years. The latter requires us to predict the future of connectivity, while the former doesn’t. And that’s exactly what *I* want to do.” To my surprise she eventually agreed, and I went on a quest to identify the future of connectivity. The title on my business card changed from “General Manager” to “Strategic Business Development” and I started my new role.

Technology trend analysis

While it sounds obvious after reading this book (remember, the book was written *after* the following took place), I decided that the first step should be to analyze technology trends that can affect the need for connectivity. While there are different technology parameters along which connectivity can improve (mobility, speed, power consumption, etc.), I chose *speed* as the parameter to follow. There are two trends that create the need for speed: *rich content*, that requires larger space to store, and *storage capacity* growth, that allows storing such rich content as well as creating new usage models. I started researching those trends.

Rich content

With time, the demand for richer content continuously grows. Digital cameras advanced from an average of 1 megapixel in 1998 to an average of 5 megapixel in 2004. There is no sign of this trend slowing down or leveling off. This requires storing and uploading bigger picture files. The average JPEG camera picture file in 1998 was 100KB. It was 4MB in 2004. Camcorders followed suit. They are capturing higher quality video, requiring more storage to hold. High-Definition MPEG2 camcorders were available, requiring more storage. Samsung demonstrated the MINIKET Flash-based camcorder, and Panasonic claimed that a 2GB Flash can hold 3 hours of MPEG4 video. At the same time, there is a transition to playing back richer content that can be taken “on the go”. DVDs were transitioning from the standard definition movies (using 5.7GB of storage) to high-definition (HD) movies (using 25GB of storage). Technologies such as Blu-Ray and HD-DVD with capacity of 25GB and 50GB became available. There was more content “on the go.” Starting with Flash-based MP3 players capable of playing 60 minutes of MP3 music—moving on to iPod with 40GB of storage (and ability to carry 10,000 MP3 songs, the transition was continuing to Portable Media Players (PMP) that had the capability of playing video movies anywhere, anytime. I estimated that the next generation was going to be capable of playing HD movies. Finally, new business models such as “local caching” of new video releases in a DVR-like device (such as Disney and Intel’s “Moviebeam”), that could be downloaded to such a portable player were driving the need to make rich content transfers faster.

Storage capacity

Between the demand for richer content and larger storage capacity, it was unclear to me which was the “chicken” and which was the “egg”. Did the tremendous growth in storage capacity lead to demanding richer content, or did the demand for richer content drive demand for higher storage capacity? One thing was clear to me—storage capacity growth outpaced any technological evolution in the past few decades, even exceeding Moore’s law.

The average hard disk drive (HDD) in 1996 was 3.5” in diameter, turned at 3,600RPM, and held 540MB of storage. The average HDD in 2004 was 2.5”, turned at 9,600RPM, and held 160GB of storage. “Business 2.0” magazine identified the capacity growth of HDD as one of the seven most significant technology trends in 2004. HDD storage was also decreasing in physical size. Samsung showed at CES 2004 products including 5GB 1” HDD, and privately have shown 0.85” drives holding 4GB, aiming at the handset market. Those HDD were what made iPod a reality, and were expected to move into many other highly portable devices, creating new applications.

Using a trend line and factoring the information above, I projected that by 2008 the ***average*** (not maximum) hard disk drive could hold 2,750GB, will be 2.1” in diameter, and will turn at 15,700RPM, maintaining the current capacity annual growth rate of 104% (double Moore’s law growth).

In 1989 Intel introduced Flash memory. The first Flash component had a capacity of 1MB. In 2004 Samsung and others introduced Flash memories with 1GB capacity—three orders of magnitude bigger. Maintaining the trend line, I projected that by 2008 there could be a 6.4GB Flash memory ***chip*** (following Moore's law closer, with annual storage capacity growth rate of 58%).

Bottlenecks

Increasing interface speed deals with "moving bottlenecks." The transfer speed along a data path is limited by the slowest link in that path. As one link in that data path becomes the bottleneck, it will typically "get upgraded" and the bottleneck will move elsewhere. Considering the path from a Flash-based digital camera into a PC HDD, the data path includes:

- Flash memory card read path;
- External cable interface (at the time—that majority used USB 2.0);
- PC internal DDR memory interface;
- PC internal DRAM write speed;
- PC internal DRAM read speed;
- ATA –type HDD interface; and—
- HDD write speed

In the following discussion, I focused on the speed required for data transfer, assuming there is no bottleneck, in order to define what is the speed required for a new, higher speed external cable interface (mostly, USB). While I needed to keep in mind that there might be other, internal bottlenecks—I assumed that those will be solved as they become the bottleneck. SD, MMC, ATA (and Serial ATA), DDR, PCI, PCI-Express and others have all increased their throughput capabilities along the years as each became a bottleneck. I therefore focused on the external cable interface bottleneck.

How fast is fast enough?

In order to determine what is the speed required for the entire data path (and thus for the external cable interface element of it) I considered three factors: the application, the forecasted capabilities of hard disk drives, and the forecasted capabilities of flash memory.

I chose a few applications that typically used (or could use) external cable interface. Those applications were measured in *capacity*, and I compared the *time* it takes to transfer a unit of data for each of those applications. I had to consider what was reasonable for the end user, using the different possible USB speeds, including what I referred to as the *SuperSpeed* interface, assuming it would run at 10Gbps).

	Song / Pic	256 Flash	1GB Flash	SD-Movie	HD-Movie	160G HDD
	4 MB	256 MB	1 GB	6 GB	25 GB	160 GB
USB 1 FS	5 sec	6 min	22 min	2.2 hr	9.3 hr	2.5 day
USB 2 HS	0.1 sec	8.5 sec	33 sec	**3.3 min**	14 min	1.5 hr
SuperSpeed	0.01 sec	0.4 sec	**1.7 sec**	**10 sec**	**42 sec**	**4.4 min**

Figure 1 – How fast is fast enough?

Note: It was assumed that the real throughput is 50% of the protocol maximum signaling rate. USB runs typically at 50% "real throughput" (e.g. the throughput of USB 2.0 was closer to 240Mbps rather than the stated 480Mbps, which is really the signaling rate).

Was the speed difference between USB 2.0 and the SuperSpeed external interface important enough for product users? The answer was that we care about speed in *some* cases, but not in *all* cases. I needed to determine whether we cared in *2008* rather than in 2005, as 2008 was when I expected the SuperSpeed interface to be launched commercially.

When do we not care?

- Backing up the 160GB PC hard disk drive over USB 2.0 took 1.5 hours. While the SuperSpeed interface could take it down to 4.4 minutes, backup is a function typically done in the background, or when the computer is not busy with anything else, and thus we might not care for the speed improvement. Of course, I had to consider the 2,750GB average HDD in 2008, which would take 36 minutes using SuperSpeed at 10Gbps compared with 12.2 hours using USB 2.0 at 480 Mbps. That's when it starts to matter.

- Transferring a *single* 5MP (megapixel) JPEG picture or a single MP3 song using USB 2.0 reduced the transfer time to 0.1 seconds from 5 seconds (one might argue that this is not long anyway) using USB 1.1. SuperSpeed will further reduce this time to 0.005 seconds, but I assessed that we would neither care nor notice that difference.

When do we care?

- Downloading a purchased movie into a portable media player's hard disk drive would take 3.3 minutes using USB 2.0, and 10 seconds using SuperSpeed. However, when it comes to a high-definition movie it would have taken 14 minutes using USB 2.0, and only 42 seconds using SuperSpeed. For the function of downloading a movie "to go"—14 minutes was unacceptable, in my opinion.

- Downloading the entire content of a Flash-based digital camera storage (1GB on average in 2005) to the PC, it could take 33 seconds (reasonable) using USB 2.0. However, relying on flash storage capacity projections of 6.4GB in 2008 (combined with growing picture quality that will result in the same number of pictures), the download is expected to take 3.5 minutes (not reasonable anymore as the user typically waits by the computer for this operation to complete), and only 11 seconds using the 10Gbps SuperSpeed interface.

Considering capacities and content richness in 2005 may have not warranted a need for a connectivity speed faster than the 480Mbps available through USB 2.0. However, it became clear that it would *not* be fast enough for the projected 2008 content and capacity. This was the first indication that a much faster interface would be required by 2008.

Hard Disk Drive transfer speed

In 2005, I forecasted 2,750GB hard disk drive capacity in 2008, or an annual capacity growth rate in capacity of 104%. However, I had to also consider the *transfer* speed increase as the one driving the need for *interface* speed, or else the transfer rate might be the actual bottleneck and not the USB interface. To calculate the expected growth in data transfer rates in hard disk drives I had to consider all the factors enabling capacity growth:

- Number of disk surfaces;
- Data/Bit density over the same surface; and—
- Size of surfaces

The number of surfaces and their diameter/area declined since 1996, so I immediately discounted that as a way to grow capacity. The only way capacity grew over the 1996 to 2005 period was through data density. Therefore, the only way to grow capacity was to grow data (bit) density. The density growth takes place in two dimensions:

- Number of tracks per surface; and—
- Number of bits per track

While the number of tracks per surface does not affect the transfer speed (as data is transferred serially from only one track at a time), the number of bits per track does. Therefore I assumed that the number of bits per track grows by the square root of the overall data density growth (and the number of tracks grow at the same rate). If disk density grows 104% annually, then overall data density grows 104% annually, and therefore the number of bits per track grows 43% annually (while number of tracks per surface grows 43%, too).

The other factor affecting the transfer speed is the disk turn speed. The disk turn rate grew from 3,600RPM at 1996 to 9,600RPM in 2004, or 13% annually. With the number of bits per track growing 43% and the turn speed growing 13% annually, I projected that the disk transfer speed will grow 6.8 times by 2008, and 2.6 times more by 2012. The common 2005, 3.5" HDD transfer rate was 800Mbps. The transfer rate for 2.5" HDD was 600Mbps. The transfer rate for 1" micro-drive (typically used for highly mobile CE devices) was 80Mbps. Based on the above analysis, I projected that the transfer rate of a 2.5" HDD will be 4.08Gbps in 2008, and the transfer rate for a 1" micro-drive will be 544Mbps in 2008 (and 1.4Gbps in 2012).

Flash memory transfer speed

I focused on the READ speed of Flash memory that will be more suitable for capture devices such as a digital camera, camcorder, or just the standard USB flash drives. In 2005, the read speed out of a Flash memory was 80MB/s (640Mbps). Even then, the speed exceeded the USB 2.0 capabilities threefold. At the time, Intel had announced the Strata™ Flash memory that was going to have a read speed of 1.73Gbps. The USB interface was clearly going to be the bottleneck in getting data out of a Flash drive.

Conclusion

Given the analysis above, I concluded that the speed of the next generation external connectivity should be no less than 2.5Gbps (assuming 50% throughput resulting in 1.25Gbps) in 2008, and closer to 10Gbps in 2012. Since I expected the lifetime of that technology to be over 4 years—I believed the next generation connection should be capable of delivering data rates of at least 5Gbps, ten times that of USB 2.0 that prevailed at the time.

Why USB?

At this point of the analysis it was still not clear to me that the SuperSpeed interface should be the next generation of the *USB* (as opposed to any other interface, in existence or that needed to be developed). This was the time to analyze whether it should be USB or any other standard. To start with, I asked why it *should* be USB (3.0), but also why not to use other technologies that could support the same data speeds. Several substitute technologies could be considered to perform the same task. The list included:

- 10GB Ethernet;
- PCI Express Generation II (or even I);
- External Serial ATA;
- CE-ATA (although at a much lower speed – CE-ATA is currently defined up to 52MB/s);
- 1394b (S3200 or beyond, although it is not defined yet); or—
- Another interface, yet to be defined.

I considered several factors: ubiquity, power-carrying, installed base, and product replacement cycle.

Ubiquity and user experience

The applications that needed the new SuperSpeed interface in most cases already existed and used a specific interface. This interface was USB 2.0. Switching from USB to another interface that has a completely different usage model could cause confusion. It was preferred to maintain the same usage model, use the same cables, same drivers, same user interface, etc. Educating users how to use a new and different interface technology always delays market adoption. The best example was the comparison of the adoption rate of USB 1 versus the adoption rate of USB 2. USB 2 was adopted dramatically faster than USB 1 because USB 1 already existed, and USB 2 was backward-compatible with it. In 2005, USB had already existed for 9 years[i].

Carrying power

One of the biggest problems with most of the alternative interfaces was the fact that they did not carry power to charge the connected peripheral device. The usage model of the USB cable allows for the power carried by the cable (from the host to the peripheral) to charge the batteries of the peripheral. I considered that a major value add for USB.

Installed base

This boils down to one major problem—the above mentioned connectors (other than USB) do not exist in PCs or peripherals today for the purpose of the desired applications. At the same time, in 2005, the market research firm InStat forecasted that by 2008, with moderate growth, over 200 million personal computers, 89 million digital cameras, 141 million flash drives, 132m personal media players, and many other devices will have USB ports. That represented a very large installed base that would warrant being compatible with. While a USB 3.0 device will work with an existing USB 2.0 port (or even USB 1.1 port)—a "new" interface will not work with any existing installed base products, and will only have value with new products “at the other end.” Using Metcalfe’s law, the value of the new interface in the new devices will be minimal until it exists on the other side. The adoption of FireWire was slow compared to the adoption of USB 2.0. The reason was that USB 2.0 built on the existence of USB 1.1 in many devices in the installed base, whereas FireWire had value driven only from new devices, equipped with FireWire in both ends.

Product replacement cycle

Another consideration was the product replacement cycle—the time we use a product from the moment we purchased it and until we buy its replacement. The replacement cycle for personal computers was relatively short (2-3 years) and thus introducing a new connector or interface technology meant that within 2-3 years the entire computer installed based will have this interface. However, it would have taken a long time until such value is built in the peripheral/consumer electronics side, which gets replaced at a much slower pace. Starting a new connectivity technology (with a new cable and a new connector) meant that for a long period of time there would have been two connectors and two cables in both sides, the PC and the peripherals. With computers and consumer electronics products becoming smaller and smaller—fewer connectors is better.

For all the reasons listed above, it became very clear to me that the SuperSpeed external cable interface will have to be the next generation USB specification—USB 3.0 (which was eventually named USB 3.0 SuperSpeed). However, given that USB is a specification created by a special interest group lead by Intel (rather than a standard created by a standard development organization such as IEEE)—it also became very clear to me that I needed to convince the powers at Intel to go down the USB 3.0 route. The participants in standards activities become a closely knit group, and in the Spring of 2005, I attended an industry meeting in Japan, and ran into Jeff Ravencraft (the chairperson of the USB promoters group) and Brad Hosler (one of the key technical talents behind USB). I asked them whether they considered going down the path of USB 3. They hesitated and told me they had thought about it, but abandoned it in favor of a newer, more efficient external cable interface such as the External Serial ATA. “Would you let me try to convince you to take the USB 3 route?” I asked. I realized that (1) the thought had crossed their mind and wasn’t completely rejected, and (2) that they could be convinced. They agreed to meet at their headquarters in Hillsboro, Oregon. We scheduled the meeting for June 2, 2005.

Taking the USB case to Intel

I had already done all the required analysis that supports the need and justification for USB 3.0 as the SuperSpeed interface, operating at a data rate between 5 and 10Gbps. All I had to do now was to put it in a presentation format, and harness all the persuasion capabilities I had.

The meeting started early in the morning. There were more than Jeff and Brad in the conference room. I pleaded the case to why USB 3.0 is the right next step in external connectivity. It seemed that the installed base, backward compatibility, and no need for a new connector were the influencing factors.

However, the Intel participants stated that they don’t have the resources to conduct a feasibility study and, frankly, didn’t believe that a 14-foot long USB cable could carry such a high data rate. “I have the resources to demonstrate that,” I stated. I lied. At the time I had no support within Texas Instruments for my USB 3 effort, and ever since I gave my business unit away—I did not have any resources I could count on. Still, they agreed to see a demo later that month.

You can get in trouble for this…

I flew back to Dallas and started seeking out the best high-speed connectivity engineers we had in our group. I explained the rationale for USB 3, and I asked for their help in developing a demonstration that will show how we can move data between two computers at a 5Gbps rate using a 14-foot cable. In full disclosure, I told the engineers who were interested in helping that (1) this is not an approved project, (2) they can get in trouble for working on it, and (3) if they do—there is nothing that I could do to bail them out, as I wasn’t the business unit general manager anymore. Little did I know (until then) that those three risks made them want to participate in this project even more. The best of them started developing a demo, and sure enough, in two-three weeks we have a crude demonstration of 5Gbps connectivity between two computers over a 14-foot cable. We scheduled the meeting at Intel for later that month (June, 2005), flew there, and the team gave the demonstration to Intel’s people. It was downhill from there, at least on the Intel front, as I still needed to convince the Texas Instruments’ management. That, on the other hand, promised to be an up-hill battle. In the first week of July, Texas Instruments and Intel exchanged a Memorandum of Understanding regarding the formation of the USB 3.0 promoters group. In general, it stated that TI and Intel will jointly promote the development of a USB Promoters Group, that both companies would be members of the USB

Promoters Group, and that the primary objective of the USB Promoters Group would be to promote the development of an Industry Standard Specification based upon the work jointly done by both companies, and that the release of the resulting specification(s) would be as timely as possible.

The Business Plan

That same month (July, 2005), I prepared a comprehensive internal business plan. The plan included not only the justification to pursue USB 3 from an industry perspective, but also how I expected this to turn into business for Texas Instruments. The business plan used the USB 2.0 adoption rate as a basis to project the adoption of USB 3.0. While the USB 2.0 specification was released in 2000, it first appeared in computers and Flash drives in 2002. Based on the growth in USB 2.0 penetration from 2002 to 2004 (the last year with empirical data prior to 2005), I estimated that in the year that the USB 3.0 specifications will be released, close to 750 million devices will have USB ports. Boy, was I wrong… While the USB 3.0 specification was released in 2008 as I expected—that year, it appeared that more than 2 billion USB-enabled devices were already shipping in 2006. Almost three times more than I estimated. I presented the business plan to my supervisor and her supervisor, a Senior Vice President at Texas Instruments. As I expected, it was challenged. Not because the numbers were not attractive—because they were. Simply because they did not believe it will happen soon enough and, more important, they wanted to increase revenue 20% every year, and cared less about increasing revenue by 100% in four years. At that time, the Connectivity group at TI merged into the ASIC group, and I decided to move on to another business unit in the company. Nevertheless, the USB 3.0 work continued at TI, albeit at a slower rate than I would have hoped for, giving a better standing to some of the company's competitors.

Bowling with a Crystal Ball

After I had moved to another business unit, I continued to track the development of USB 3.0 inside of Texas Instruments, although not as closely as before. As a Director of Strategy and Industry Relations to the Mobile Connectivity Solutions at the time, I had more time to think about what had happened, and decided to put my experience down in writing. In June of 2006, I began writing this book. I started by researching technology trends beyond those that led to the creation of USB 3.0. Not only was I interested in many different technology trends—I also wanted to understand *why*. Why were those trends so predictable? I conducted the interviews described in Part 1 of this book, conducted primary research, and finished writing in May of 2007. I did originally leave a placeholder for this chapter, due to the importance of showing how everything in this book worked. However, in May 2007 Jeff was still not sure whether Intel was going to announce USB 3.0 that year or the following year. As a result, I decided to publish the book without that chapter. The first edition of the book was published in July 2007. The USB 3.0 Promoter Group was announced two months later…

USB 3.0 is announced and released

The formation of the USB 3.0 promoters' group was announced by Pat Gelsinger, then the Chief Technology Officer for Intel, during his keynote address at the Intel Developers Forum in San Francisco on September 18, 2007. He announced the core promoters group, made of HP, Intel, Microsoft, NEC, NXP, and *Texas Instruments*. This was the first time Texas Instruments was invited to be part of the USB promoters group. He promised release of the specifications in 2008, and that it would be 10 times faster than USB 2.0, provide better power efficiency, and also mentioned that over 6.2 billion USB devices had shipped since 2001, of which 2.1 billion shipped in 2006 alone[ii]. The specification itself was released in November of 2008, as promised. Although my *technical* contribution was minimal, my name is listed on these specifications as a contributor.

USB 3.1 (SuperSpeed+)

In January 2013, the USB Promoter Group announced their intention of releasing a USB 3.1 SuperSpeed+ specification, doubling the speed of USB 3.0, reaching speeds of 10Gbps[iii], as proposed in my original 2005 business plan. This specification only took six months to develop, and by the end of July 2013, the USB Promoter Group announced the release of the specification[iv]. Synopsis, a leading IP solutions provider for Systems on Chip (SoC) was the first to market with a design that was first to market with a solution[v].

Summary

The story of creation of USB 3.0 was not only the culmination of the work in this book—it actually led to the writing of this book. This story shows the importance of accurately forecasting technology trends, understanding the opportunities resulting from those trends and their relevance to any business, and how to navigate them to fruition through the complex maze of the industry ecosystem.

---

[i] USB 1 (low speed and full speed of up to 12Mbps) was released in 1996. USB 2.0 (at 480Mbps) was released in 2000.
[ii] *USB 3.0 to arrive at superspeed in 2008*, http://www.bit-tech.net/news/hardware/2007/09/19/usb_3_to_arrive_at_superspeed_in_2008/1
[iii] *SuperSpeed USB (USB 3.0) Performance to Double with New Capabilities* (http://www.usb.org/press/USB-IF_Press_Releases/SuperSpeed_10Gbps_USBIF_Final.pdf)
[iv] *USB 3.0 Promoter Group announced availability of the USB 3.1 Specification to increase SuperSpeed USB to 10 Gbps* (http://www.usb.org/press/USB-IF_Press_Releases/SuperSpeedUSB_10Gbps_Available_20130731.pdf)
[v] *Synopsys Demonstrates Industry's First SuperSpeed USB 10 Gbps Platform-to-Platform Host-Device IP Data Transfer* (http://news.synopsys.com/2013-12-10-Synopsys-Demonstrates-Industrys-First-SuperSpeed-USB-10-Gbps-Platform-to-Platform-Host-Device-IP-Data-Transfer)