Evolving Standards:
Moore's Law and Deep Packet Capture Storage Models

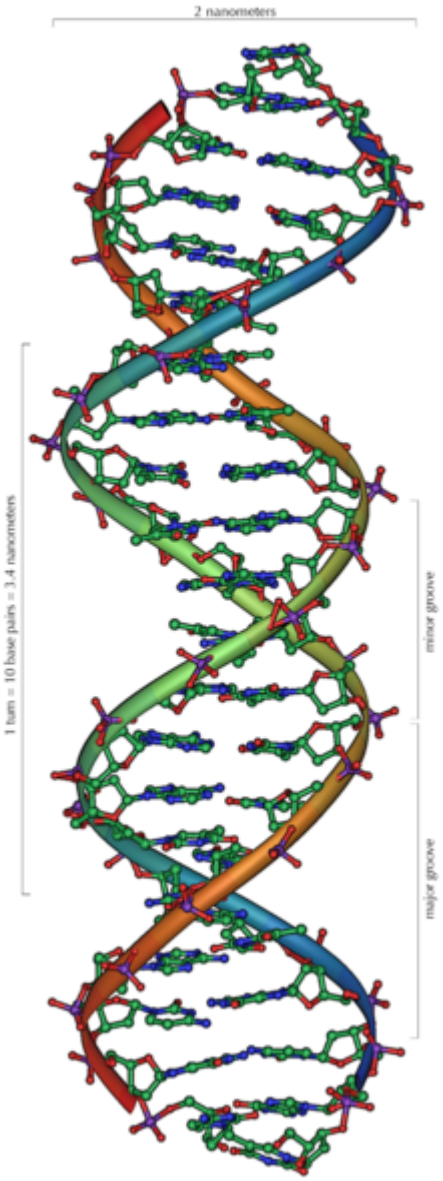

Deoxyribonucleic acid (*DNA*), the organic program that spawned the first known computer – the human mind.
This program has been running continuously for over 4 billion years.

Moore's law describes a long-term trend in the history of computing hardware. Since the invention of the integrated circuit in 1958, the number of transistors that can be placed inexpensively on an integrated circuit has increased exponentially, doubling approximately every two years, and resulting in the cost factors being reduced by roughly half. The trend was first observed by Intel co-founder Gordon E. Moore in a 1965 paper. It has continued for almost half of a century and is not expected to stop for another

decade at least and perhaps much longer.

Almost every measure of the capabilities of digital electronic devices is linked to Moore's law: processing speed, memory capacity, even the number and size of pixels in digital cameras. All of these are improving at (roughly) exponential rates as well. This has dramatically increased the *usefulness* of digital electronics in nearly every segment of the world economy. Moore's law describes this driving force of technological and social change in the late 20th and early 21st centuries.

On 13 April 2005, Gordon Moore stated in an interview that the law cannot be sustained indefinitely: "It can't continue forever. The nature of exponentials is that you push them out and eventually disaster happens" and noted that transistors would eventually reach the limits of miniaturization at atomic levels.

Extrapolation partly based on Moore's law has led futurists such as Vernor Vinge, Bruce Sterling, and Ray Kurzweil to speculate about a technological singularity. Kurzweil projects that a continuation of Moore's law until 2019 will result in transistor features just a few atoms in width. Although this means that the strategy of ever finer photolithography will have run its course, he speculates that this does not mean the end of Moore's law:

> Moore's law of Integrated Circuits was not the first, but the fifth paradigm to forecast accelerating price-performance ratios. Computing devices have been consistently multiplying in power (per unit of time) from the mechanical calculating devices used in the 1890 U.S. Census, to [Newman's] relay-based "[Heath] Robinson" machine that cracked the Nazi [Lorenz cipher], to the CBS vacuum tube computer that predicted the election of Eisenhower, to the transistor-based machines used in the first space launches, to the integrated-circuit-based personal computer.

Thus, Kurzweil conjectures that it is likely that some new type of technology (possibly optical or quantum computers) will replace current integrated-circuit technology, and that Moore's Law will hold true long after 2020. He believes that the exponential growth of Moore's law will continue beyond the use of integrated circuits into technologies that will lead to the technological singularity. The Law of Accelerating Returns described by Ray Kurzweil has in many ways altered the public's perception of Moore's Law. It is a common (but mistaken) belief that Moore's Law makes predictions regarding all forms of technology, when it actually only concerns semiconductor circuits. Many futurists still use the term "Moore's law" in this broader sense to describe ideas like those put forth by Kurzweil.

The technological singularity is a theoretical future point of unprecedented technological progress, caused in part by the ability of machines to improve themselves using artificial intelligence.

Statistician I. J. Good first wrote of an "intelligence explosion", suggesting that if machines could even slightly surpass human intellect, they could improve their own designs in ways unforeseen by their designers, and thus recursively augment themselves into far greater intelligences. The first such improvements might be small, but as the machine became more intelligent it would become better at becoming more intelligent, which could lead to an exponential and quite sudden growth in intelligence.

Others, most prominently Ray Kurzweil, define the singularity as a period of extremely rapid

technological progress. Kurzweil argues such an event is implied by a long-term pattern of accelerating change that generalizes Moore's Law to technologies predating the integrated circuit and which he argues will continue to other technologies not yet invented. Critics of Kurzweil's interpretation consider it an example of static analysis, citing particular failures of the predictions of Moore's Law.

Robin Hanson proposes that multiple "singularities" have occurred throughout history, dramatically affecting the growth rate of the economy. Like the agricultural and industrial revolutions of the past, the technological singularity would increase economic growth between 60 and 250 times. An innovation that allowed for replacement of virtually all human labor could trigger this singularity.

Good (1965) speculated on the consequencesof machines smarter than humans:

> "... Let an ultraintelligent machine be defined as a machine that can far surpass all the intellectual activities of any man however clever. Since the design of machines is one of these intellectual activities, an ultraintelligent machine could design even better machines; there would then unquestionably be an 'intelligence explosion,' and the intelligence of man would be left far behind. Thus the first ultraintelligent machine is the last invention that man need ever make..."

Communications Advancements:

Moore's Law and Speed of Light Limitations

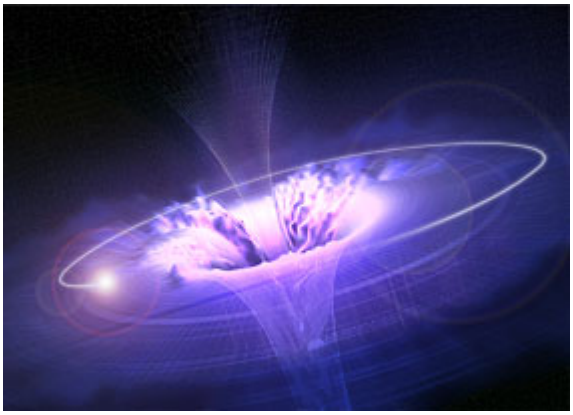

Event Horizon of a Black Hole from which even light cannot escape.

In February of 2008, the NSF requested US $20 million from the U.S. government for fiscal 2009 to start the "Science and Engineering Beyond Moore's Law" effort, which would fund academic research on technologies, including carbon nanotubes, quantum computing and massively multicore computers, that could improve and replace current transistor technology.

A quantum computer is a device for computation that makes direct use of quantum mechanical

phenomena, such as superposition and entanglement, to perform operations on data. The basic principle behind quantum computation is that quantum properties can be used to represent data and perform operations on these data.

Although quantum computing is still in its infancy, experiments have been carried out in which quantum computational operations were executed on a very small number of qubits (**qu**antum **bi**nary digi**ts**). Both practical and theoretical research continues with interest, and many national government and military funding agencies support quantum computing research to develop quantum computers for both civilian and national security purposes, such as cryptanalysis.

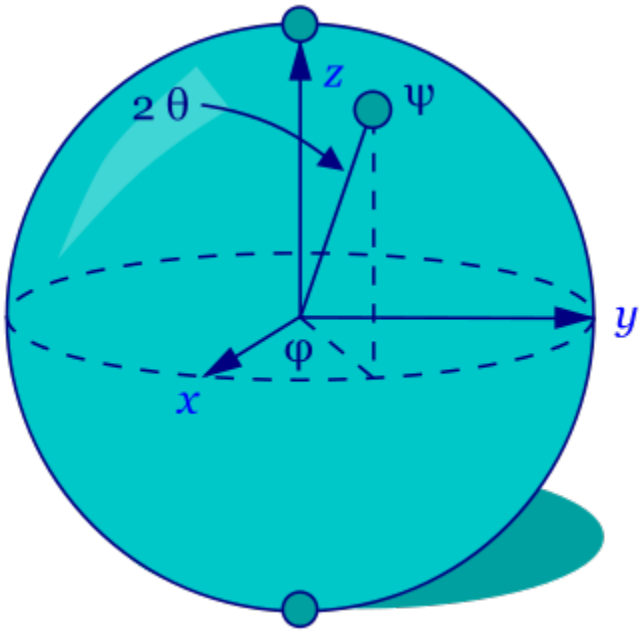

The Bloch sphere is a representation of a qubit,

the fundamental building block of quantum computers.

If large-scale quantum computers can be built, they will be able to solve certain problems much faster than any of our current classical computers (for example Shor's algorithm). Quantum computers are different from other computers such as DNA computers and traditional computers based on transistors. Some computing architectures such as optical computers may use classical superposition of electromagnetic waves. Without some specifically quantum mechanical resources such as entanglement, it is conjectured that an exponential advantage over classical computers is not possible.

Network Storage and communications models are strongly affected by Moore's Law, since they for the most part are completely tied to the underlying computer hardware and its communications capabilities. Networking Technology also exhibits the exponential growth rates in capabilities and performance levels seen in current semiconductor evolution and it logically follows these trends will also continue into the realms of quantum computing, since conceivably, even these systems will require communications capabilities.

It is inevitable that Networking trends and raw data delivery capabilities will continue to dramatically

advance over the next 20 years, reaching perhaps into the Terabyte or Pedabyte per second ranges. However, there are theoretical limits in physics dealing with analog and digital signaling technologies when applied to speed of light restrictions, even with circuits based on Gallium Arsenide (GaAs) and other alternative materials which exhibit dramatically reduced signal propagation delays.

In quantum mechanics, certain quantum effects may appear to be transmitted at speeds greater than the speed of light. For example, the quantum states of two particles can be entangled. Until the particles are observed, they exist in a superposition of two quantum states. If the particles are separated and one of them is observed to determine its quantum state, then the quantum state of the second particle is determined automatically.

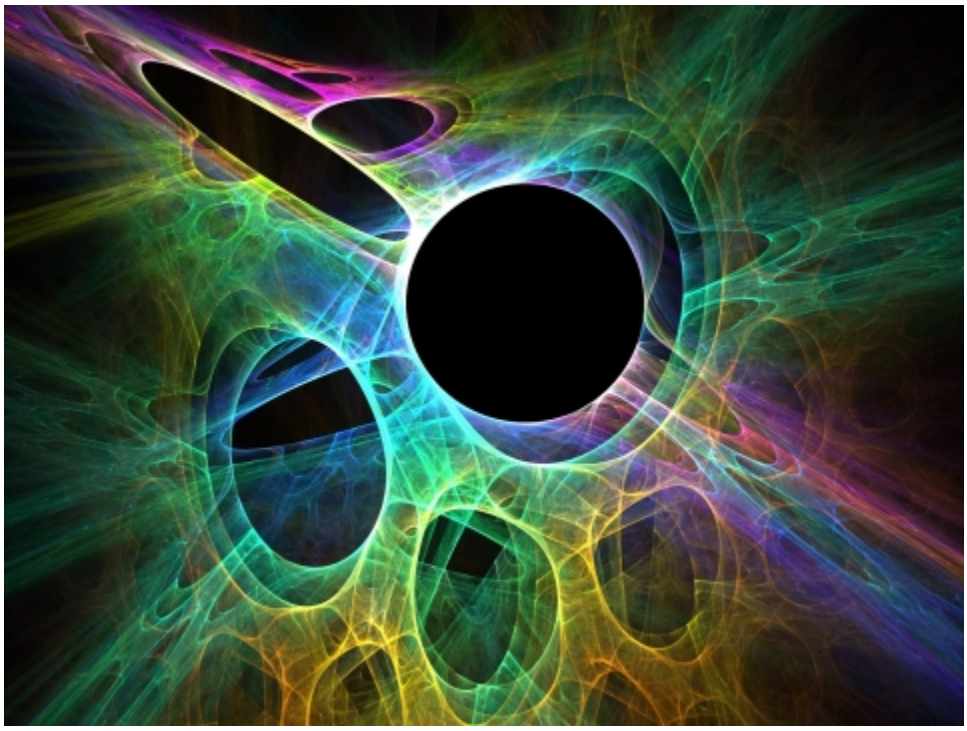

Subatomic particles at the quantum level postulated by the super-symmetric string theory

Quantum entanglement is a quantum mechanical phenomenon in which the quantum states of two or more objects are linked together so that one object can no longer be adequately described without full mention of its counterpart — even though the individual objects may be spatially separated. This interconnection leads to correlations between observable physical properties of remote systems. For example, quantum mechanics holds that states such as spin are indeterminate until such time as some physical intervention is made to measure the spin of the object in question. It is equally as likely that any given particle will be observed to be spin-up as that it will be spin-down. Measuring any number of particles will result in an unpredictable series of measures that will tend more and more closely to half up and half down.

However, if this experiment is done with entangled particles the results are quite different. When two

members of an entangled pair are measured, one will always be spin-up and the other will be spin-down. The distance between the two particles is irrelevant.

In order to explain this result, some have theorized that there are hidden variables that account for the spin of each particle, and that these hidden variables aredetermined when the entangled pair is created. It may appear then that the hidden variables must be in communication outside of space-time no matter how far apart the particles are, that the hidden variabledescribing one particle must be able to change instantly when the other is measured. If the hidden variables were to suddenly stop interacting when they are far apart, the statistics of multiple measurements must obey an inequality (called Bell's inequality), which is, however, not what is observed -- both by quantum mechanical theory, and in actual experiments and observation.

Quantum teleportation, or entanglement-assisted teleportation, is a technique used to transferinformation on a quantum level, usually from one particle (or series of particles) to another particle (or series of particles) in another location via quantum entanglement. It does not transport energy or matter, nor does it allow communication of information at superluminal (faster than light) speed. Its distinguishing feature is that it can transmit the information present in a quantum superposition, useful for quantum communication and computation.

It is possible for shock waves to be formed with electromagnetic radiation which appear to bend space-time as a local effect, irregardless of the reference frame for any given unit of time. If a charged particle travels through an insulating medium faster than the speed of light in that medium then radiation is emitted which is analogous to a sonic boom and is known as Čerenkov radiation.

Čerenkov radiation (also spelled Cerenkov or Cherenkov) is electromagnetic radiation emitted when a charged particle (such as an electron) passes through an insulator at a speed greater than the speed of light in that medium. The characteristic "blue glow" of nuclear reactors is due to Čerenkov radiation. It is named after Russian scientist Pavel Alekseyevich Čerenkov, the 1958 Nobel Prize winner who was the first to characterize it rigorously.

It is generally considered that it is impossible for any information propagated as electromagnetic radiation or as matter to travel faster than the speed of light , because it would travel backwards in time relative to some observers. However, there are many physical situations in which speeds greater than the speed of light are encountered. Some of these situations involve entities that actually travel faster than the speed of light in a particular reference frame, but when measured with the reference frame of the observer, none of the information or matter appears to travel faster than the speed of light as an observable electromagnetic phenomena.

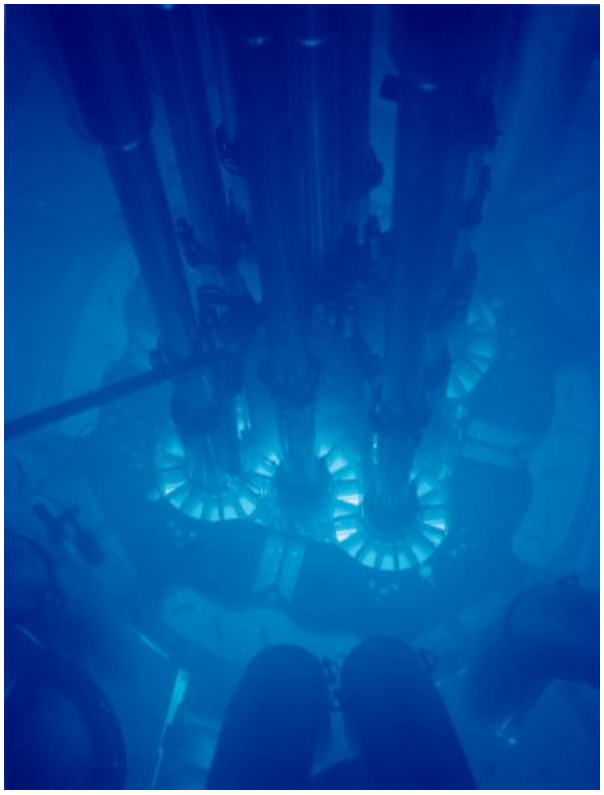

Čerenkov Radiation Effect observed as visible light from the

Advanced Test Reactor located at the Idaho National Laboratory

It is possible for the "group velocity" of light to exceed the speed of light and in an experiment in 2000, laser beams traveled for extremely short distances through caesium atoms with a group velocity of 300 times the speed of light. It is not, however, possible to use this technique to transfer information faster than the speed of light since the velocity of information transfer depends on the front velocity, which is always less than the speed of light.

Exceeding the group velocity of light in this manner is comparable to exceeding the speed of sound by arranging people distantly spaced in a line, and asking them all to shout "I'm here!", one after another with short intervals, each one timing it by looking at their own wristwatch so they don't have to wait until they hear the previous person shouting. Another example can be seen when watching ocean waves washing up on shore. With a narrow enough angle between the wave and the shoreline, the breakers travel along the waves' length much faster than the waves' movement inland.

If a laser is swept across a distant object, the spot of light can easily be made to move at a speed greater than the speed of light. Similarly, a shadow projected onto a distant object can be made to move faster than the speed of light. In neither case does any matter or information travel faster than light relative the current observers frame of time when observed as a purely electromagnetic phenomena.

Despite the restrictions imposed by speed of light, the relationship of Moore's Law to the trends in communications technology can be expected to continue, since irregardless of speed of light limitations, the ability to communicate information through quantum teleportation would facilitate a technological singularity in the field of communications. Although the propagation of information through space time

would still be limited by speed of light constraints, the cross sectional bandwidth of a quantum based communications fabric could be incomprehensibly large and could be created "at will" in any width or size deemed necessary by a quantum computer, and discarded when not in use. The real challenge would be more geared towards making such technology practical for intergalactic or interplanetary communications across light years using quantum fabrics to provide ISP capability to a distant colony given the limitations imposed by light speed. These restraints, however, would not be appreciably measurable on networks connected between New York and Los Angeles with a quantum fabric.

Quantum computers could conceivably create at will vast and on demand communications capabilities, since information could be transferred through quantum teleportation. Speed of light is only a meaningful constant in terms of its use to describe a much simpler relationship between the physical laws that define the universe – its nothing more than a heuristic ratio based upon observations of the nature of the Universe itself. The universe and its laws "bend" in relationship to the absolute ratio express in our comprehension of the relationship we call "the speed of light". In computer science terms, the universe and the interactions between its simple laws emerge as a *synchronous* event model.

Deep Packet Capture: History

HP 4952A / Agilent 4952A SDLC Protocol Analyzer with 18180 Interface (RS-232C /V.24/RS449)

Arguably, some of the first true deep packet capture applications were the HDLC/SDLC analysis probes employed in early 3270 development by IBM and later ATT for use in Mainframe communications development efforts. They were large and expensive hardware boxes used to snoop on SNRM frames to debug communication line problems with SCC based HDLC/SDLC terminal controllers. 802.2 LLC and Token Ring came next, and were the dominant networking technologies for large systems until IBM and Novell made Local Area Networking ubiquitous and pervasive. The internet came much later, of course.

Network General set the standard for what later became deep packet capture, but little changed in several decades in this technology area – sniffers and deep packet capture were relegated to "development and troubleshooting tools" status, and in many ways are still trapped in this model. Their storage models at the time being little more than collections of simple MSDOS files.

Network analysis trace files have taken different forms over the years, and additional changes have been made as new analyzer versions and topologies were created. For a large portion of the early history of Local Area Network adoption in corporate America, the Network General Sniffer format was considered

the common format in the network analysis industry. Most network analysis tools could save their trace files into the 'Sniffer' format for compatibility between analyzers.

In the original Sniffer, Ethernet trace files were saved as a format ending with the extension .ENC. Token ring traces had the extension .TRC, FDDI had .FDC, etc. These files contained a frame-by-frame output of the captured data in their original raw hexadecimal format.

Network General Compaq 3 based Sniffer (mid-1990s)

As technology progressed, the size of captured trace files became larger and larger. Network General changed their format to include the option to compress the file whensaving. Although this created a much smaller capture trace on disk, the file extension didn't change. This meant that trace files saved with this default compression could not be loaded into other analysis tools. Many manufacturers and developers have reverse-engineered the compression methods used by Network Generaland can load compressed Sniffer trace files. A large number of sniffer "knock-off" products such as Novell's LANAlyzer appeared as software only solutions to compete against the Sniffer, however, most of them failed to achieve any significant traction because at that time, hardware and communications platform capabilities had not co-evolved to make software only solutions in deep packet capture viable.

When the Windows version of SnifferPro was released, the file formats changed. The .xNC extensions were replaced with a different format ending in .CAP. Not only did this format change completely, all topologies used the .CAP. Because of this generic naming convention, it became necessary to save any Sniffer Pro trace files into a subdirectory that described the topology, or a name for the file with the

topology included.

The .CAP files are non-compressed trace files. A documented-but-rarely-used feature of Sniffer Pro is to save the files with the .CAZ extension. Any files saved with a .CAZ will automatically be compressed by Sniffer Pro.

Some applications do not save their output files into Sniffer or Sniffer Pro format, such as Microsoft's Network Monitor. To help with this conversion, WildPackets' ProConvert utility can change almost any network analysis format into another.

libpcap was originally developed by the tcpdump developers in the Network Research Group at Lawrence Berkeley Laboratory. The low-level packet capture, capture file reading, and capture file writing code of tcpdump was extracted and made into a library, with which tcpdump was linked. It is now developed by the same tcpdump.org group that develops tcpdump.

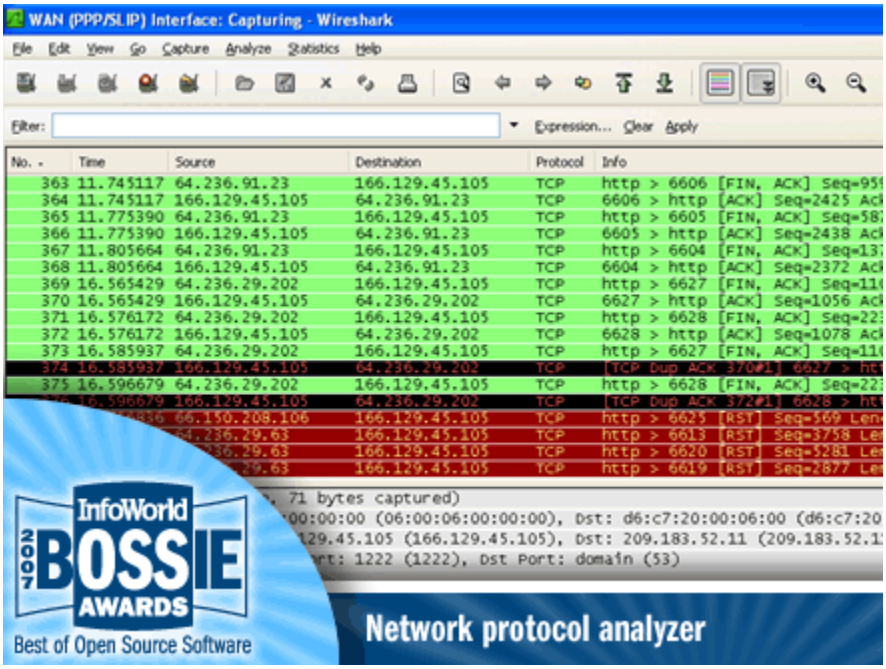

Wireshark has all but replaced Network General's Sniffer with a software-only product released under an open source license.

LIBPCAP has become ubiquitous and is now deployed as a standard for virtually every packet capture and network analysis tool in use today, displacing CAP/CAZ as the format of choice for long term storage of captured packet data.

Software models have invariably always been the beneficiaries of Moore's Law, since any increase in computing power or bandwidth can be immediately realized by software independent of the underlying hardware platform which hosts it. Software evolution and its relationship to hardware closely parallels the disposable soma theory of biology, since any software program views the hardware that serves as its container as entirely disposable – any many programs such as Internet worms express this characteristic in the highest degree being completely mobile and self-replicating. Arguably, some of the most

advanced distributed software in existence today are the botnets that exist across the Internet. Many of these malicious applications rival even the most sophisticated distributed computing models created for legitimate commercial applications. It can be said with certainty that commodity platforms are the end result of Moore's law – cheap and ubiquitous computing power for the masses.

During the late 1990's, the ability to create sustained deep packet capture and stream to disk capability at the then 100Mbit speeds was not possible without custom hardware, as the platforms of the day were cost prohibitive and Moore's Law had not reached the threshold of bus and disk speeds necessary to sustain these rates with commodity hardware, which spawned an entire industry of ISP based adapters and vertical hardware solutions to offload packet processing and classification from the processor and I/O bus. However, close analysis of the development of commodity stream to disk hardware shows an interesting parallel. 1Gb Ethernet emerged in the commodity market beginning in 1997 with commercially viable offerings, however, it was five years until commodity platforms had the horsepower to actually build and sustain software only solutions to enable deep packet capture, Niksun and Network General being the first companies to actually build multi-gigabit stream to disk solutions, with Solera Networks being the first to provide these capabilities in an open platform. The same trend has followed with 10Gb Ethernet. When 10Gb first appeared, PCI Express and multi-lane controllers were not available to sustain the rates necessary to stream 1.2 gigabytes/second to disk. Current platforms are now rapidly approaching the ability to sustain these rates, though the cycles between hardware advances appears to be seven years between the emergence of 10Gb and the availability of commodity platforms to enable this ability with software only solutions.

Immediately after a new communications threshold attains commodity status, such as the emergence of 10Gb ethernet, the trend has been for hardware vendors to deploy custom capture and analysis cards for deep packet capture while the speeds and feeds of the underlying platforms slowly catch up to the performance levels of vertical adapter solutions, ultimately rendering these same hardware solutions obsolete. This trend with the emergence and die off of hardware adapter-assisted deep packet capture solutions has been repeated over and over again and is now into the fourth iteration over a twenty year span with first 10Mb Ethernet, then 100Mb Ethernet, 1Gb Ethernet, and now 10Gb Ethernet. Just like a thunderstorm in the desert can transform a dried up desert arroyo into a lush garden for a week until scorched and withered by the sun when the water evaporates, Deep Packet Capture hardware adapter solutions are boom or bust for the first few years after the emergence of a new commodity network standard, then slowly die off while software solutions and faster platforms erode away the vertical vendors dependent on proprietary hardware and their market share, and then the cycle starts again.

Free software's rise is linked more to the beneficial effects of Moore's law on advancing software capabilities than the proposition that “free is better.” Free software became viable when commodity platforms became powerful enough to run a Unix system on a home computer. However, in a world of quantum computing, much of what Unix is today would necessarily become irrelevant and would be discarded in favor of self defining Operating Systems capable of creating and self modifying their own code at the quantum level. Moore's law defines that technology leaps every two years, which means what people are using today will probably be in a landfill 4 years from now.

Deep Packet Capture: Convergence

The list of deep packet capture vendors and solutions spans from simple programs that write capture files on traditional file systems to complex vertical and proprietary storage solutions. Many of the deep packet capture solutions are just more of the same seen at various intervals and are tied to the "rest" cycles of Moore's Law which naturally follow the release of a new communications threshold: specialized adapters to take up the failings in current platform performance and disk speeds, vertical storage solutions, and vertical applications directed at various classes of networking customers.

Virtually all of the vendors in the landscape are targeted at vertical applications for network performance, monitoring, analysis, reconstruction, and searching of captured data, and have the same basic elements in their vertical solutions. The solutions provided by these vendors are also becoming more similar to one another and are converging. One indicator of this is the reliance on a common storage architecture for exporting and importing data. There are now over a dozen companies creating storage devices and support LIBPCAP as a common standard, and almost all of the Deep Packet Capture vendors can import and export LIBPCAP data formats with their technology.

Deep Packet capture devices are also no longer simply "development tools", and many of these platforms are becoming active systems for not only monitoring the network, but actively managing it as well. Convergence of Network Routing and Deep Packet Capture Platforms are increasingly integrated with Intrusion Detection Systems that control network firewalls with human intervention occurring after the fact when alert logs and reviewed, and these solutions are becoming more and more software based. Hardware based Deep Packet Capture solutions and their boom and bust cycles are logically tied to the "rest" cycles following the release of a new commodity communications threshold, after which software solutions slowly catch up as they benefit from the incremental effects of Moore's Law in the underlying hardware platforms.

Moore's Law is perhaps the most pervasive influence driving the creation and adoption of standards in the computer industry. As commodity systems flood the market, standards are created not by prescience, but based upon actual needs. It can be said that the one nice thing about standards is there are so many of them to choose from – and most of them are inevitably abandoned as the market shifts from one paradigm to another as Moore's Law leaps the industry forward. Customers and people in general vote with their feet more so than their hands and the technologies deployed and used on commodity systems dictates the framework for standards. In the absence of a standards body, vendors create technology solutions based on market need and customers vote for the standards based upon their deployment and adoption. Very few standards bodies have been prescient enough to anticipate every customer requirement or trend in the technology sector. In fact, the vast body of evidence indicates that most standards or proposed standards from the 1990s long ago fell out of use or were replaced, and not surprisingly, as a result of advancements predicted by Moore's Law. IPX/SPX was the standard for local area networking through the 1980s and 1990s, but when the Internet became accessible to the general public it quickly disappeared and fell out of favor and IP replaced it – based on need.

Network General's sniffer formats became a standard for deep packet capture over a decade based on need, not elegance or sophistication, it became a standard by default – simply because there was nothing else to do the job and a lack of public will to drive another standard – and the needs were minimal.

The landscape of Deep Packet Capture has dramatically altered in the past 5 years. Before the advent of stream to disk capabilities, sniffers and vertical solutions required complex triggers and distributed event monitoring to identify what happened on a Network for timing related problems or software defects. This drove the need for more and more complex sniffer logic and custom hardware adapters since these devices were never designed to capture all of the traffic and payload traversing a network, and could not review the events that occurred on the network except real time. The end result was an entire industry of products and solutions based upon this event driven model, and the advent of stream to disk capability created a paradigm shift that left Network General and other vendors scrambling when Niksun began marketing one of the earliest stream to disk appliances. Although Niksun's offerings were quite humble at the time, they were a significant technology leap over the more traditional sniffer approaches to network management and troubleshooting.

Now there are dozens of vendors producing stream to disk appliances for network archival, driven for the most part by competitive pressures and legislative compliance for many of their customers. As the Internet and networking has become an integral element woven into the fabric of modern society, it has driven the need for standards. Stream to disk and deep packet capture storage is poised to enter the mainstream market as commodity technology at the storage level, since post capture analysis coupled with real time network monitoring is the real value proposition Deep Packet Capture Storage provides with the more advanced stream to disk technologies.

Moore's Law has produced commodity servers that can sustain 10Gb stream to disk speeds. A common public storage architecture standard for multi-gigabit packet capture archival and retrieval for forensics and mining applications is going to be chosen by the masses, and if no one is talking about it, then they will do what they have always done, vote with their feet, and the commodity offerings will unquestionably gain the greatest penetration. Packet capture storage is going to be a commodity in the next 24 months – Moore's Law is alive and well.

Jeffrey Vernon Merkey

Forensic Filesystems

www.forensic-filesystem.com

About the Author: Mr. Merkey is a former senior computer scientist of Network Associates and is the former Chief Scientist of Solera Networks. Mr. Merkey is also a former Chief Scientist of Novell. Mr. Merkey is the creator of Network General's Infinistream platform and the Solera Networks DS Series Appliances, and is the inventor of four patents in the field of Networking, Storage, and Operating Systems.

Much of the reference material for this article was used from research articles on Wikipedia and other creative commons sources, along with Mr. Merkey's opinions related to network forensics and deep packet capture.