December 7, 2017

TECHNOLOGY/SEMICONDUCTORS & COMPONENTS

Semiconductors: Technology and Market Primer 10.0

SUMMARY

The Oppenheimer *Semiconductors: Technology and Market Primer 10.0* provides a comprehensive overview of the semiconductor industry, both from a technology and market perspective.

KEY POINTS

- Our *Semiconductor Primer 10.0* is an updated version of the "one-stop-shop" resource for the dynamic and complex semiconductor industry. This tenth edition is designed for investors new to the space as well as for those seeking a better understanding of key technological and market elements.
- Inside you will find plain-English discussions of the technology behind the semiconductor industry and its major end markets. We discuss the major drivers of sector fundamentals, its cyclicality, and emerging trends. Finally, we dive into the individual product and end market segments, including forecasts and market share for both semiconductor vendors and their hardware customers.

Rick Schafer
720-554-1119
Rick.Schafer@opco.com

Joshua Buchalter, CFA
212-667-8387
Joshua.Buchalter@opco.com

Disseminated: December 7, 2017 23:00 EST; Produced: December 7, 2017 16:05 EST

Introduction

The Oppenheimer *Semiconductors: Technology and Market Primer 10.0* is an updated version of the comprehensive "one-stop-shop" resource for the dynamic and complex semiconductor industry we first issued in June 2003 and updated in October 2004, December 2005, January 2007, January 2008, December 2009, October 2011, October 2013, and November 2015.

The report is targeted to those investors new to the sector as well as those looking for a comprehensive resource to help them better understand key technological or market elements within the industry. We also suggest it as a "desk reference" for more experienced investors, as it has lots of forecast and market share data as well as in-depth discussions of many of the important trends affecting the semiconductor industry.

We start with basic semiconductor definitions and a simple review of manufacturing processes. We then discuss the semiconductor cycle and key fundamental indicators, and introduce some important concepts. We follow with a discussion of the major semiconductor product groups, including revenue forecasts and market share data. We then take a deep dive into the most important end markets served by the semiconductor industry, highlighting key players and company market shares for each application.

Product summaries cover analog, microcomponents, logic, memory and discrete devices. End market summaries cover computing, networking, telecom/datacom, wireless, digital consumer and automotive semiconductors. We've also included a section on "emerging technologies," where we briefly discuss technologies that are just getting off the ground and that will be exciting to watch in 2017 and beyond.

Any comments that will help us make this a more successful document are welcomed.

Rick Schafer
Managing Director

Joshua Buchalter
Associate

Oppenheimer & Co.
Semiconductor Equity Research

Note on this version: All historical data on companies is current as of 2Q17; and industry sales, unit, and utilization data is also current up to January 2017. Market share data has been updated to reflect the most recent third-party sources and our internal estimates. All forecasts have been updated as well, and were also extended to 2021 (from 9.0's 2019). Naturally, there were also changes to formatting and the report structure to make it more readable and relevant.

Acknowledgments

We would like to acknowledge the contributions of our research colleagues at Oppenheimer whose help was invaluable in the preparation of this report.

Communication Technology

Ittai Kidron
George Iwanyc
Michael Leonard

Emerging Technologies and Services

Andrew Uerkwitz
Martin Yang
Paul Dean

Enterprise/Applications Software

Brian Schwartz
Koji Ikeda
Tyler Page

Communications, Infrastructure Software and Israeli Technology

Shaul Eyal
Hugh Cunningham
Tanner Hoban

Communications Services

Timothy Horan
Tom Shaughnessy
Matthew Wilson

Internet

Jason Helfstein
Jed Kelly
Alec Brondolo

Table of Contents

Section I: Semiconductor Basics

This section deals with the basics of semiconductors

Topics include:

Semiconductor definitions

Manufacturing

- Semiconductor manufacturing process
- CMOS
- Lithography & Moore's Law
- Wafer Size
- Manufacturing Strategies
- Geographic Trends
- Semiconductor Capital Equipment

Semiconductor Definitions

A **semiconductor** is a solid-state substance that is halfway between a conductor and an insulator. When charged, the substance becomes conductive; when the charge is eliminated, it loses its conductive status.

By combining conductive material, semiconductor material, and insulators in a pre-determined pattern, the movement of electricity can be precisely controlled. Semiconductors are therefore ideal for building devices that control the operation of electronic equipment.

A **transistor** is the basic element used in building semiconductor devices. A transistor is fashioned from semiconductor material and acts as an on/off switch, which opens and closes when electrically activated.

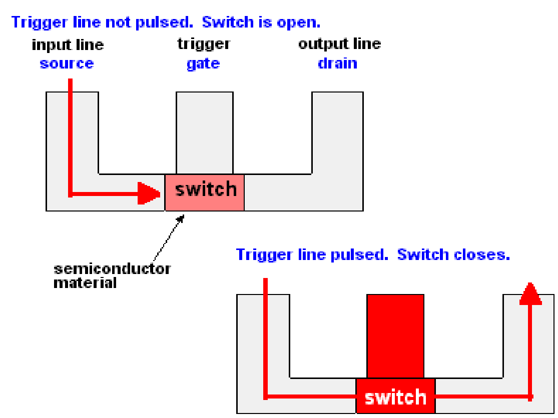

Source: Computer Desktop Encyclopedia, Oppenheimer & Co.

A semiconductor is a solid-state substance that is halfway between a conductor (which conducts electricity very well) and an insulator (which doesn't conduct electricity at all). When charged with electricity or light, the substance becomes conductive, allowing electricity to flow through it. When that charge is eliminated, it loses its conductive status, and electricity cannot flow through. By combining conductive material, semiconductor material, and insulators in a pre-determined pattern, the movement of electricity can be precisely controlled. Semiconductors are therefore ideal for building devices that control the operation of electronic equipment.

The basic element used in constructing semiconductor devices is the transistor. A transistor is essentially a tiny on/off switch fashioned from semiconductor material, which sits between two charged regions known as a "source" and a "drain." The switch can open and close when electrically activated, allowing current to flow from source to drain (when the switch is closed or "on"), or blocking the current's passage (when the switch is open or "off"). A third region, the "trigger gate," controls whether the switch is open or closed. By manipulating the gates within a semiconductor device, the system can accurately control the movement of electricity.

Semiconductor Definitions

The simplest semiconductor devices are comprised of a single transistor; these are called **discretes** and are usually used to control the flow of signals and power within a larger electronic device.

More complex semiconductor devices are built by combining multiple transistors and conductive **interconnect** material to form **logic gates**. These logic gates are arranged in a pre-defined pattern to perform more complex processing or storage functions. These devices are called **integrated circuits (ICs)**. The most integrated of these devices are known as **system-on-a-chip devices**, or **SoCs**.

The simplest semiconductor devices are comprised of a single transistor. These devices are referred to as discretes, and they are used in all types of electronic equipment to control the flow of signals and power within a larger electronic system.

When many transistors are combined, an integrated circuit is created that can be used to process or store data signals in an electrical format. Engineers design pathways using transistors and conductive interconnect material set in logical arrangements (called "logic gates") to perform specific functions, and these pathways become the circuit architecture etched onto the surface of the chip. Modern semiconductor devices can perform a variety of different functions on a single chip, and many electronic devices are powered by highly integrated system-on-a-chip (SoC) devices.

All electronic devices are built around semiconductors—transistors make up logic gates, logic gates make up circuits and circuits make up electrical systems. Key end applications are computing, telecommunications, data networking and storage, wireless communications, consumer electronics, automobiles, industrial equipment, and aerospace/military.

Semiconductor Device Structure

Semiconductor devices contain a **transistor array**, which is etched onto a rectangular piece of **silicon** called a **die**. The die is encased in a plastic or ceramic **package**, and tiny wires called **wire bonding** are used to connect the input/output gates on the chip to the leads on the outside of the package.

Typically, the cost split for a fully packaged IC is roughly 85% for the silicon and 15% for the package.

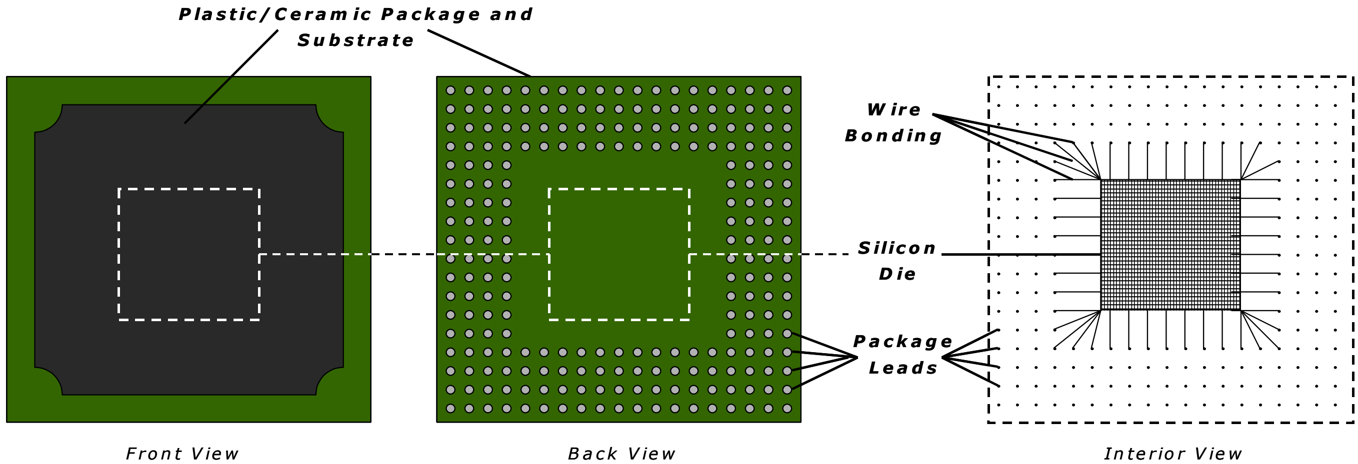

Source: Oppenheimer & Co.

Typically, semiconductor devices are sold as standalone, fully functional packaged products that are ready to be implemented in an electrical system. The transistor array (or in the case of a discrete, the lone transistor), which is etched onto a rectangular piece of silicon (called a "die") during the wafer fabrication process, is housed in a plastic or ceramic package. Tiny wires called "wire bonding" are used to connect the input/output gates on the die to the leads on the package. The semiconductor device is snapped or soldered onto a printed circuit board, with the leads attached to conductive pathways along the board.

For most semiconductor devices, the split in manufacturing value for a fully packaged semiconductor device is roughly 85% for the silicon and 15% for the package, although this can vary with the type of package and complexity of the device. Lower end devices such as discretes can have 30% or more of its value in the package given the maturity of the transistor process used at the trailing edge. On the other end of the spectrum, high-end devices such as microprocessors sometimes need more complex packages to deal with the speed and heat dissipation of the device, and therefore derive a higher value from the package as well.

Note that in many cases, IC vendors offer the same die in a variety of packages to accommodate different feature sets, power requirements, device characteristics, platforms or customers. Depending on the design, the type of package used can have a dramatic impact on the speed, power consumption, heat dissipation and footprint of the device. Some devices will even be "pin-for-pin" compatible with devices from competitors, usually at the request of the customer.

Semiconductor Devices In Systems

Getting Designed In

Semiconductor devices are used as components in larger electrical systems designed by an **original equipment manufacturer (OEM)**. Chip vendors **sample** their devices to OEMs, and are awarded **design wins** as the OEM designs the part into their systems. The chip vendor then provides production versions of the device, which are **qualified** by the OEM before going to full production.

Semiconductor devices can be custom designed for a specific customer and platform; these are called **ASICs** and are usually designed in cooperation with the OEM. Merchant devices called **ASSPs** are not designed for specific customers and can therefore be used by multiple OEMs. For the most generic components, OEMs will often buy **off-the-shelf** components, many of which are sold through the distribution channel.

Semiconductor devices are used as components in larger electrical systems designed by an original equipment manufacturer (OEM). Chip vendors supply samples of their devices to OEMs, sometimes based on specifications dictated by the customer. They are then awarded design wins as the OEM designs the part into their systems. The chip vendor will then provide production versions of the device, which will be qualified by the OEM before going to full production. This entire process can be as short as several weeks or as long as a year or even more.

Semiconductor devices can be custom designed for a specific customer and platform; these are called application-specific integrated circuits (ASICs) and are usually designed in cooperation with the OEM. ASICs are generally used when performance is key (e.g., a core router) or when volume on a single device is significant enough (e.g., a game console). Merchant devices, called application-specific standard products (ASSPs), or simply “standard products,” are not designed for specific customers and can therefore be designed in by multiple OEMs. ASSPs are generally used when time-to-market and cost are key (e.g., PCs, set-top-boxes). For the most generic components, OEMs will often buy off-the-shelf components, many of which are sold through the distribution channel. This includes not only memory but also standard logic, analog and discrete devices.

Customer engagements take many forms. Sometimes the OEM issues straight purchase orders; other times they will have a deeper engagement with a formal supply contract or even a technology partnership. The deeper the engagement, the more customization the OEM will typically expect from the chip designer. Note that design wins are awarded not only based on the performance of the device; factors such as pricing, software, service, support, existing relationship, track record, strength of the roadmap, scalability and other factors are often just as important.

Semiconductor Devices In Systems

System IP and Standards

For more complex electronic systems, semiconductor vendors will incorporate system **IP (intellectual property)** directly into their ICs. Often, they will try to integrate as many logic functions as possible into a single **SoC (system-on-a-chip)** or **chip set** in order to reduce the OEM's **bill of materials**. The IC vendor's system knowledge will be critical in winning the design.

Chip vendors design their devices to conform to the needs of their target customer set. Sometimes, the OEM will supply specifications directly to their chip vendors. In mature, high-volume markets, the devices will often adhere to specifications as defined by a **standard**, set forth by a **standards body**. This allows multiple chip and equipment vendors to compete more easily, speeding time-to-market and lowering the cost of the technology implementation.

When dealing with more complex systems, semiconductor vendors will incorporate system intellectual property (IP) directly into the device, either in the form of logic gates or in software or firmware that runs on top of the device. This is especially true in the case of application specific standard products, which get sold into multiple platforms at multiple vendors. By incorporating the system IP, the chip vendor lowers the design cost for the OEM and also speeds time-to-market, two factors that often matter more than simple performance or device pricing.

IC designers will usually try to integrate as many logic functions as possible into a single IC in order to reduce the OEM's bill of materials. Sometimes, it will prove too difficult or too costly to integrate certain functions, and the IC vendor will offer a chip set, either internally or with a partner. In any case, the IC vendor's system knowledge will be critical to winning the design.

Chip vendors will design their devices to conform to the needs of their target customer set. Sometimes, the OEM will supply specifications directly to its chip vendors, either in an ASIC arrangement or when multiple vendors compete for a standard product win. In the more mature, high-volume markets, devices will often adhere to specifications as defined by a standard, set by a standards body such as International Organization for Standardization (ISO) or the Institute of Electrical and Electronics Engineers (IEEE). This allows multiple vendors to compete, speeds time-to-market, and lowers the cost of the technology implementation across the supply chain. Chip vendors maintain seats on the standards bodies alongside their OEM customers in order to influence the outcome of standards negotiations.

Semiconductor Devices In Systems

The Evolving Electronics Supply Chain

The electronic equipment industry supply chain continues to evolve, as OEMs increasingly outsource aspects of the production and procurement process. **Electronics manufacturing services (EMS)** companies build products for their OEM customers. **Original device manufacturers (ODM)** companies go a step further, taking over portions of the design process. **Distributors** sit both between OEMs and end customers and between component suppliers and OEMs to smooth the supply chain.

Semiconductor vendors must maintain relationships with their OEM customers' outsourcing partners. Component decisions are increasingly being pushed toward the ODMs and EMS providers, favoring companies with strong supply chain relationships.

The electronic equipment industry supply chain has undergone some significant changes over the last 10-15 years. Enterprise and service provider equipment OEMs have shifted their business models away from chasing hardware margins and increasingly toward cultivating software and service businesses. Consumer device OEMs have focused on speeding time-to-market and on new channels to reach customers. Those that maintain large hardware operations are increasingly focused on design and marketing activities as opposed to hardware manufacturing and assembly. The trends have driven hardware prices lower and increased focus on costs.

In addition to shifting production to low-cost regions like Asia, OEMs are increasingly looking to outsource some aspects of the production process. Electronics manufacturing services (EMS) companies are now responsible for a big piece of PC, wireline, and wireless equipment production. Original device manufacturer (ODM) companies take it a step further, taking over some aspects of the design process. The PC industry also has a large and mature motherboard sector that serves both PC OEMs and the distribution channel, while notebooks are predominantly assembled by third parties. Distributors sit both between OEMs and end customers and between component suppliers and OEMs to smooth the supply chain and to simplify export processes in an increasingly global electronics industry.

In order to keep pace with the evolving supply chain, semiconductor vendors must maintain relationships with their OEM customers' outsourcing partners. Component decisions are increasingly being pushed toward the ODMs and EMS, favoring semiconductor suppliers with strong channel relationships. In the PC market, for example, servicing the motherboard makers is just as important as winning designs with OEMs back in the U.S and Japan. In the handset market, ODMs are just as important as OEMs. In the analog market, having the right distributor relationships plays a tremendous role in who wins and who loses.

In the diagram below, we display the different roles played by different classes of companies in delivering a semiconductor device into the hands of the end customer. Note that we display the primary function in black, with some additional functions that they can also perform in gray. We follow the diagram with a description of each class.

Electronics Equipment Industry Supply Chain Diagram

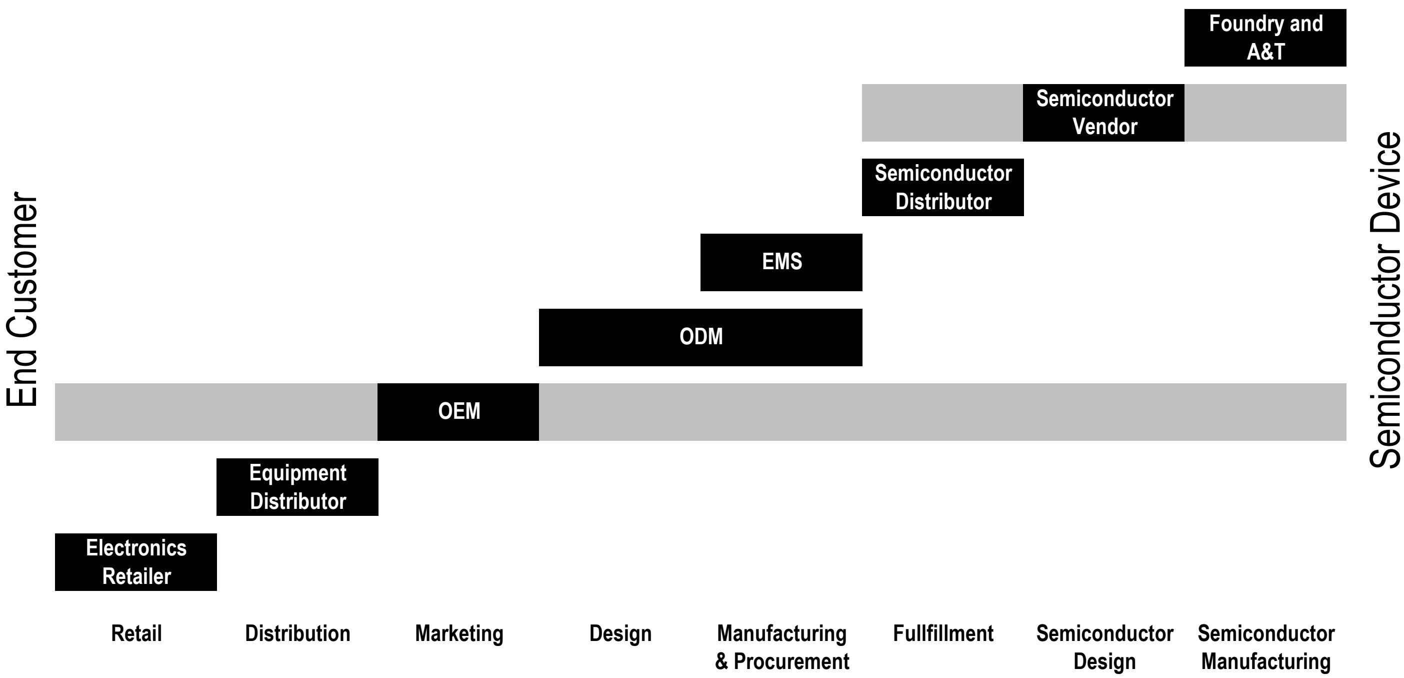

Source: Oppenheimer & Co.

- **Foundry and Assembly and Test:** These companies perform manufacturing, assembly, and testing of semiconductor devices. They are discussed in greater depth later in this report. Examples include TSMC, UMC, ASE, GlobalFoundries, SMIC and Amkor.
- **Semiconductor Vendor:** Semiconductor vendors perform chip design and marketing. Some vendors perform their own manufacturing, others use foundries and assembly and testers, and still others use a mix. On the other side, semiconductor vendors can either sell their parts directly to OEMs, ODMs, or EMS, or they can use a distributor (in practice, most use both). Examples include Intel, Texas Instruments, Broadcom, and Maxim.
- **Semiconductor Distributor:** These distributors perform three important functions: 1) they carry inventory to help smooth the supply chain, 2) they handle import logistics to simplify international shipments, and 3) they reach smaller customers that the chip vendor cannot service directly. Examples include Avnet and Arrow.
- **Electronics Manufacturing Services (EMS):** EMS companies perform manufacturing on behalf of OEMs. In some cases they can handle procurement of components as well. Examples include Flextronics and Jabil.
- **Original Device Manufacturer (ODM):** ODMs are similar to EMS companies except they go a step further, taking over some aspects of the design and procurement process. ODMs partner with OEMs or service providers, who perform marketing and distribution. Examples include Hon Hai (Foxconn), Compal, and Quanta.
- **Original Equipment Manufacturer (OEM):** OEMs are the electronics providers who design and market branded products to end customers and service providers. OEMs can pursue a wide variety of vertical integration models; some (like IBM or Fujitsu) do it all, from chip design and manufacturing all the way to direct sales to customers. Others focus on marketing and some aspects of design, but outsource or partner for most other functions. Examples include HP, Dell, Apple, Samsung, Cisco, Nokia, and Sony.
- **Equipment Distributor:** These distributors help OEMs reach customers, servicing individual retailers or chains as well as selling directly to customers (through catalog or Internet). Examples include CDW and Ingram Micro.
- **Electronics Retailer:** Retailers provide outlets for individual consumers to buy electronics. Examples include Best Buy, Walmart, and increasingly more online channels such as Amazon, Tiger Direct, and New Egg.

Semiconductor Manufacturing

Semiconductor devices are manufactured in specialized factories called **wafer fabs** using a process known as **wafer fabrication**. Circular wafers of silicon are put through a cycle of chemical processes in order to etch an ion-charged transistor array as patterned on a set of **masks**. On top of the transistor array, layers of metal **interconnect** form pathways between the transistors; the layers are insulated by a **dieletric** material.

After wafer processing, the finished wafer is put through a **dicing** process, where individual **die** are separated. These are sent to a **back end** facility for **packaging and assembly** and final **test**.

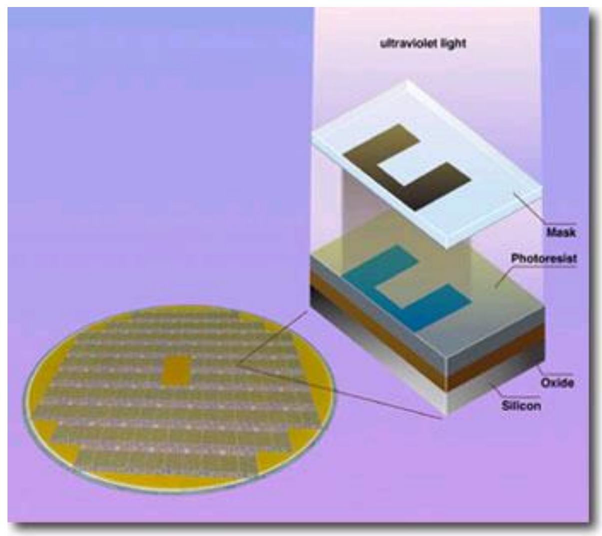

Source: IBM

Semiconductor devices begin the manufacturing process in specialized factories called “wafer fabs” using a process known as “wafer processing,” “wafer fabrication,” or “front end” manufacturing. In this set of processes, circular wafers of silicon are put through a cycle of chemical processes in order to etch a transistor array pattern on the wafer. During this process, a stepper will image the circuit pattern from a set of masks that contain the device design. Through stages of deposition, masking, etching and implantation using advanced photolithography, a series of intricate transistor arrays are formed on the surface of the wafer (multiple units are etched onto each wafer; they are later separated into individual devices).

Once the transistor array is complete, the interconnect layers are built on top of the silicon transistors. These, too, are patterned on a set of masks and are created through multiple cycles of deposition, masking and etching using photolithography. Interconnect is applied layer by layer, with an insulating dielectric material deposited in between. The interconnect material is usually aluminum or copper, while a variety of non-conductive materials are used for the dielectric.

Once wafer processing is complete, the wafer is sliced into individual die with a diamond drill through a process called dicing. The individual die is then transferred to a “back end” facility for final packaging and assembly, followed by a final test cycle.

Semiconductor Manufacturing

Semiconductor Manufacturing Flow Diagram

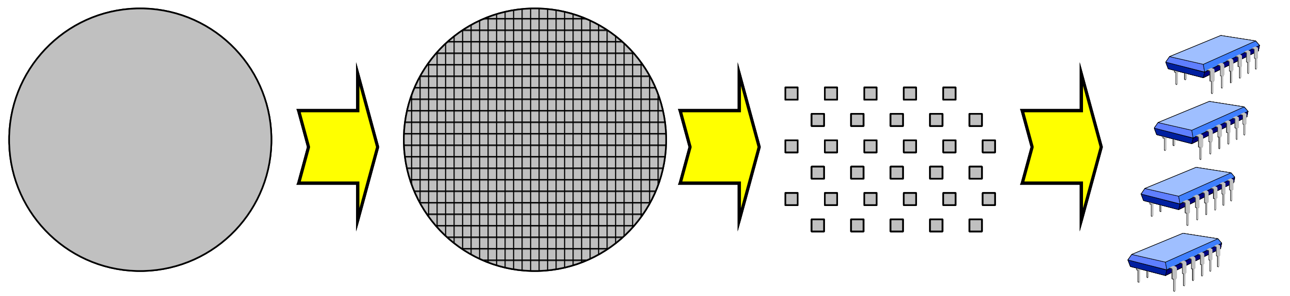

Wafer Manufacturing	Wafer Processing	Dicing	Assembly & Packaging
Key Steps:	**Key Steps:**	**Key Steps:**	**Key Steps:**
Crystal Pulling	Chemical Vapor Deposition	Dicing	Die Attach
Crystal Slicing	Coating		Wire Bond
Wafer Lapping and Polishing	Align and Expose		Encapsulation and Finishing
Cleaning	Develop and Bake		Final Test
Oxidation	Etch and Strip Resist		
	Ion Implantation		
	Physical Vapor Deposition		
	Chemical Mechanical Planarization (CMP)		
	Passivation		
	Wafer Probe/Test		

Source: Oppenheimer & Co.

Semiconductor manufacturing can be divided into four relatively discrete stages. These include wafer production, wafer processing, dicing, and back end assembly and test.

- **Wafer production:** In this stage, silicon wafers are created from raw silicon. Separate wafer companies perform this process; few semiconductor device makers produce their own wafers.
- **Wafer processing:** Wafers are run through a series of chemical and lithographical processes to etch the transistor array and interconnects. Also called front end processing, this is the longest, most complex and most costly stage of semiconductor manufacturing. “Fabbed” semiconductor producers and foundry suppliers perform this process.
- **Dicing:** The processed wafer is chopped into individual die using a diamond drill. This is done at the beginning of the back end process.
- **Assembly and test:** Individual die are placed into a plastic or ceramic package and tiny wire bonding is used to connect the input/output gates on the chip to the leads on the outside of the package. The finished device is then tested. Assembly is done either by fabbed semiconductor producers or by third-party outsourcers; test is usually done in-house by both fabbed and fabless chip producers but can be outsourced.

The steps are discussed in further detail on pages 15-19.

Wafer Manufacturing

Silicon wafers are produced by heating a mixture of silica and carbon in a furnace, creating wafer-grade silicon. A seed is then dipped into the molten silicon and is slowly twisted and pulled out. This creates a cylindrical **ingot** several feet long, which is ground to an appropriate diameter (200mm, 300mm, etc). The ingot is then sliced into thin wafers for shipment to IDMs and foundries.

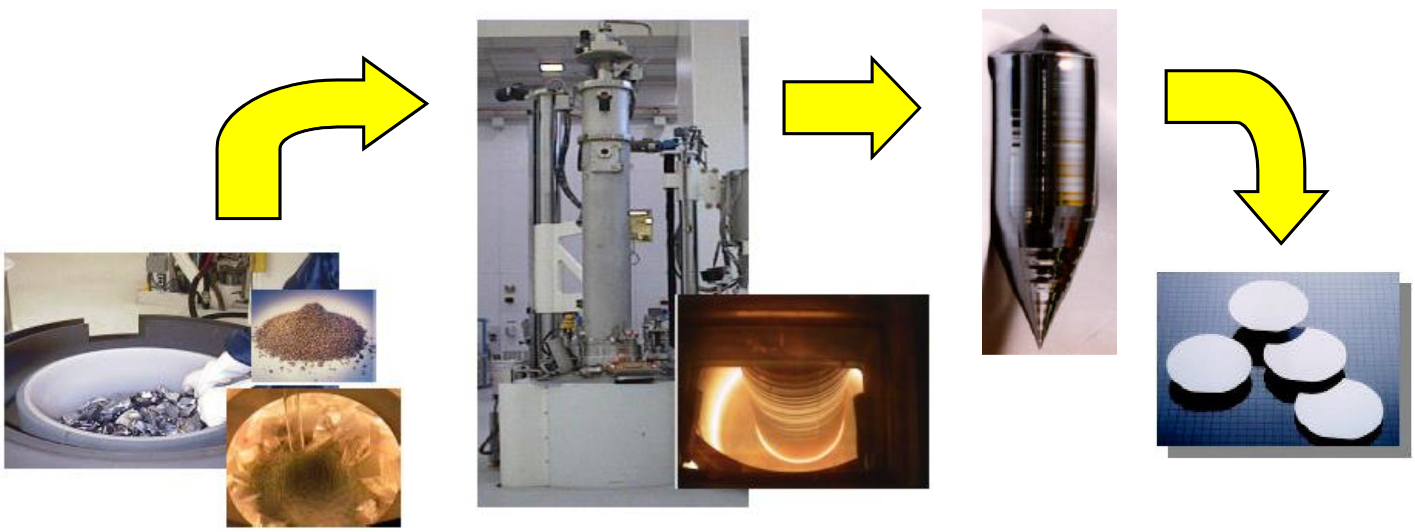

Source: MEMC, Oppenheimer & Co.

The first step in building a semiconductor device is the manufacture of silicon wafers. A mixture of silica and carbon is heated in a furnace, creating a molten mixture of wafer-grade silicon. A silicon seed is dipped into the melt and slowly pulled out. This creates a cylindrical ingot several feet long, which is ground to the proper diameter (200 mm, 300 mm, etc.). The ingot is then sliced into thin wafers for shipment to IDMs (integrated device manufacturers) and foundries. The pictures below display crystals, ingots and finished wafers.

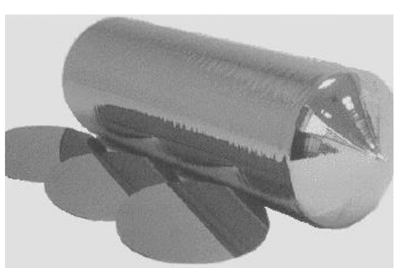

Source (all pictures): MEMC, IBM, Texas Instruments, Computer Desktop Encyclopedia

Wafer Processing: Pre-Metal

In the pre-metal stage, wafers are put through an intense cycle of chemical processes in order to etch a transistor array patterned on a set of **masks.**

Through cycles of **deposition/oxidation, photolithography, etching**, and **ion implantation**, the transistor array is created. Charged ions are then deposited to enable transistor functionality.

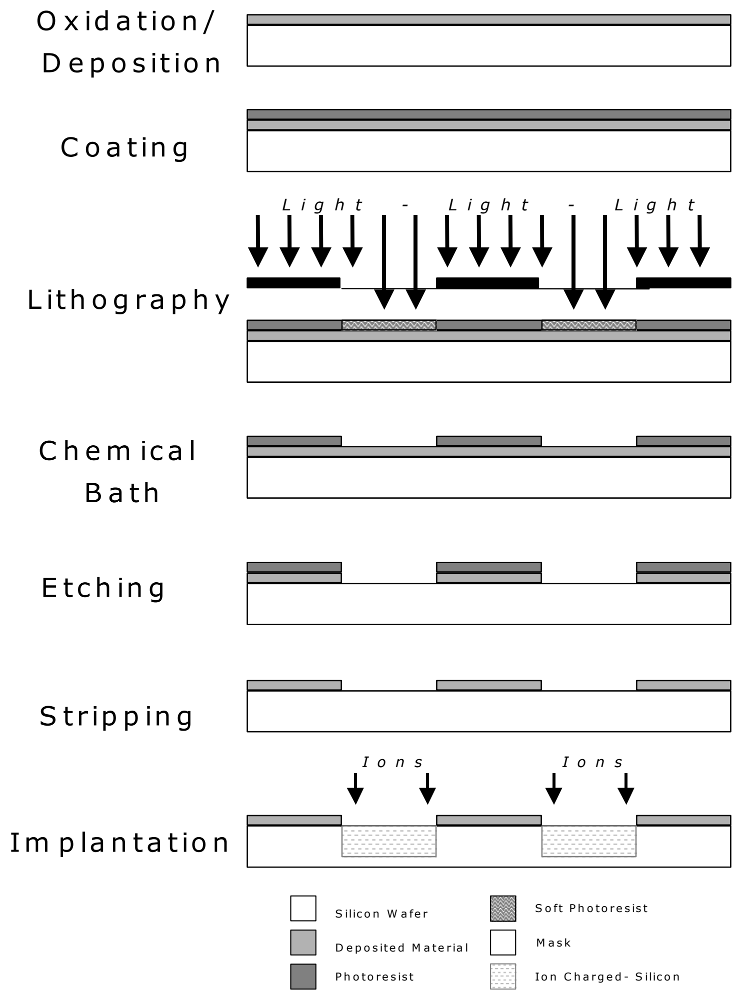

Source: Oppenheimer & Co.

The first half of the front end process is called the pre-metal stage. In this stage the transistor array is etched onto the wafer and ions are deposited in selected areas beneath the surface of the silicon.

Oxidation/deposition: In this first step, oxide material is deposited onto the surface of the wafer. This can be done through oxidation, which involves the heating of the wafer in a furnace filled with oxide gas, or through physical (PVD) or chemical (CVD) deposition of the material onto the surface of the wafer.

Photolithography: The wafer is then coated with a layer of photoresist—a material that is sensitive to light. A stepper then focuses an intense beam of light through a pre-formatted mask, softening the photoresist material in certain areas of the wafer according to the pattern of the mask in front of the light source. The wafer is then sent through a chemical bath to dissolve the soft photoresist.

Etching: Etching tools then remove the oxide material that is not still covered by photoresist. Once etched, the remaining photoresist material is stripped away using special chemicals.

Ion implantation: At various stages throughout the process, implantation equipment creates charged regions within the silicon wafer to enable transistor functionality. Ionic materials known as dopants are shot into the wafer, where they remain under the surface of the silicon. Both positive and negative dopants are used.

Wafer Processing: Interconnect

After the transistors are created, they are connected together to form logic gates using an **interconnect** material, usually aluminum or copper. The metal layers are built with stages of deposition, lithography, and etching, similar to the pre-metal stage but using separate equipment and masks. Dielectric material is deposited between the layers to insulate them from one another.

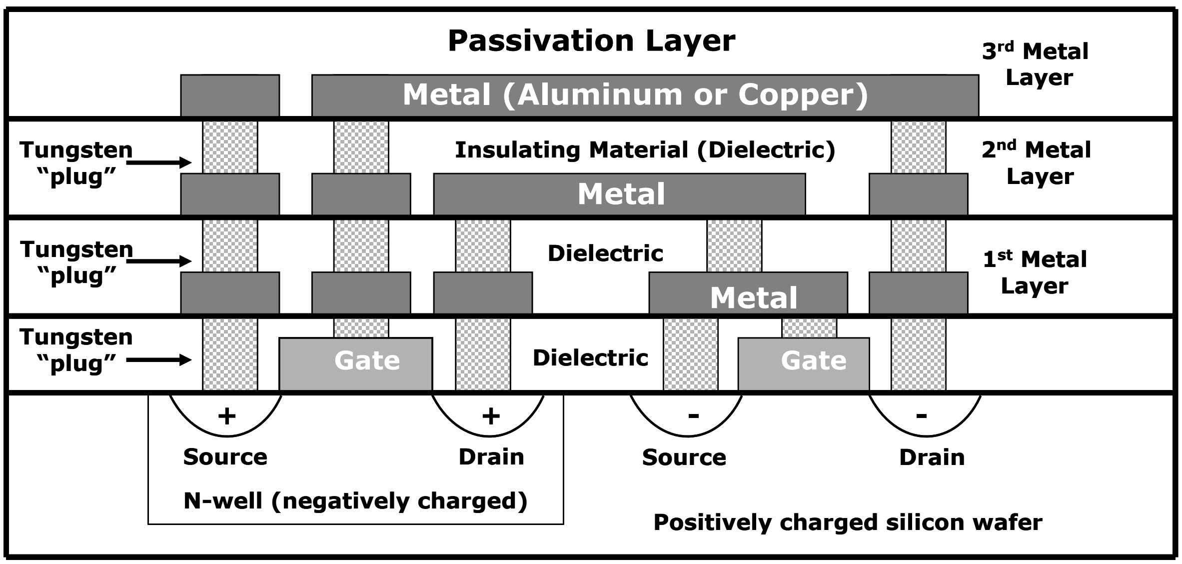

Source: Oppenheimer & Co.

After the transistor array is created, the wafer is moved to the metal area of the fab, where interconnect layers are deposited and etched to create the logic gates. The metal interconnect layers are created through cycles of deposition, photolithography and etching, similar to the pre-metal stage but with separate equipment. In between each of the metal layers, dielectric material is deposited to insulate the metal layers from one another. A final passivation layer is added above the top metal layer to protect the circuitry.

Most devices use aluminum or copper for the interconnect layers. Aluminum tends to resist the flow of electricity as wires are made thinner and narrower; therefore, most device manufacturers shifted to higher performance copper interconnects when they began 0.13-micron production. The number of metal layers varies by device; the highest performance devices can have as many as nine metal layers.

The picture below displays an IBM-fabricated SRAM cell with the insulating oxide films removed.

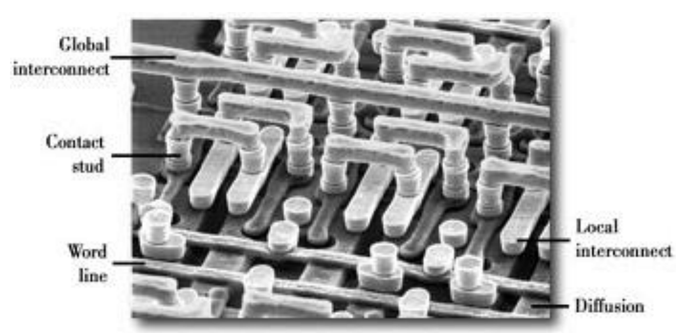

Source: IBM

Dicing

After the wafer is processed, a diamond drill is used to slice the wafer into individual **die**. Each die is tested before being sent to the back end facility.

Semiconductor Dicing Diamond Drill

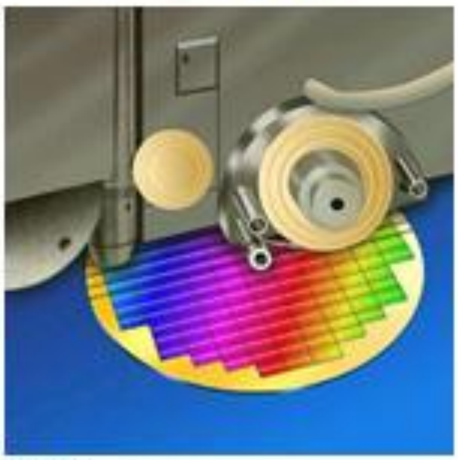

Source: IBM

Finished Wafer and Individual Die

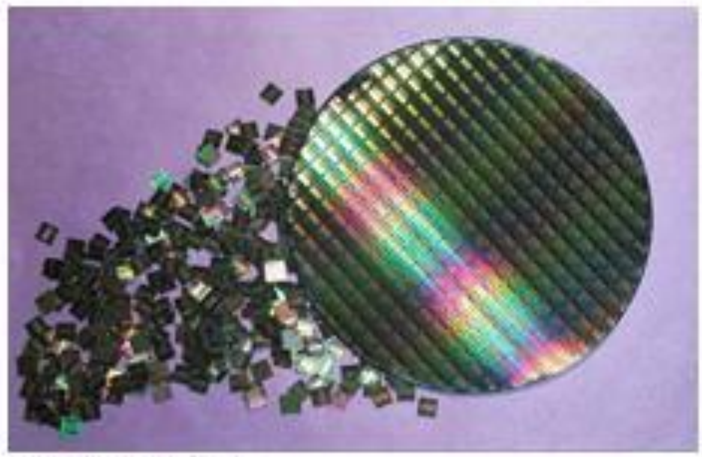

Once the wafer is fully processed and all pre-metal, metal and inter-metal layers are built, a diamond circular saw is used to separate the individual die. The die are then tested for defects and the good die are sent to a back-end line for packaging and final test.

Back End: Assembly & Test

The individual die are sent to a back end assembly facility, where they are **attached** to a package, **wire bonded**, **encapsulated**, and run through **final test**.

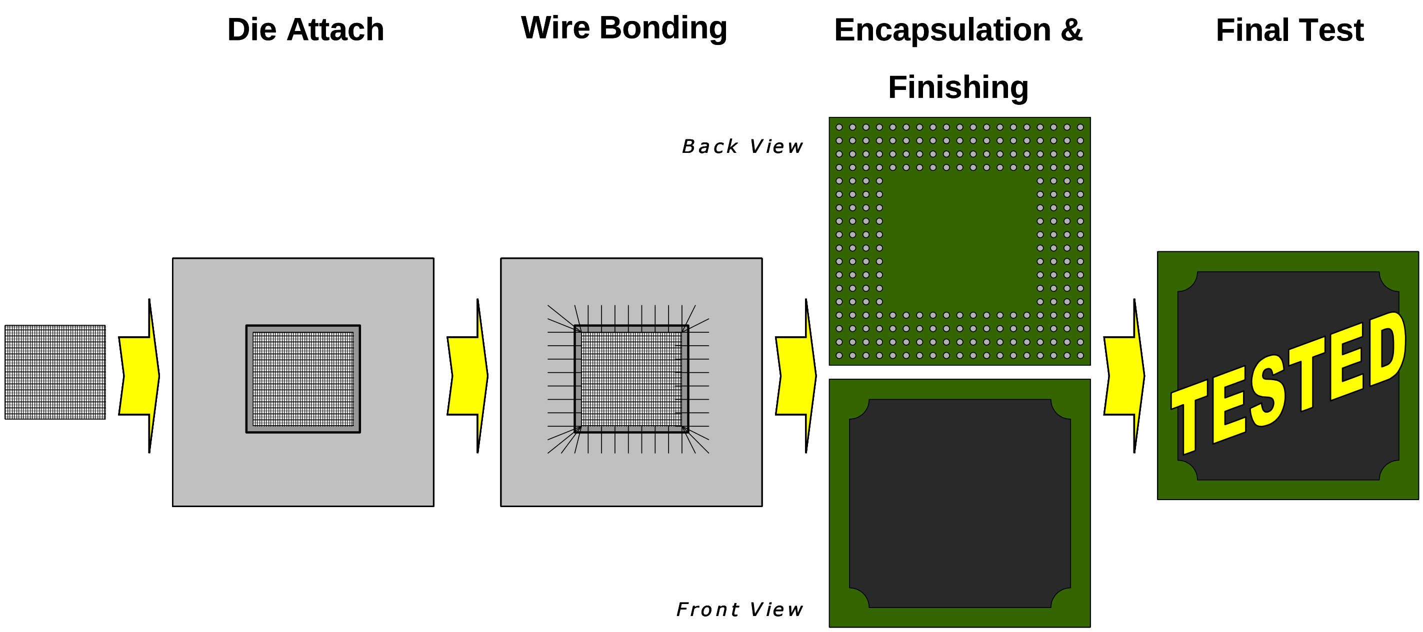

Source: Oppenheimer & Co.

Once diced, the individual die are then packaged and tested. Each die is first attached to a package, and wire bonding is used to connect the input/output gates on the die to the leads on the outside of the package. The device is then encapsulated, sealed and markings are affixed. A final test is conducted before shipment.

Note that many devices are now using a process called flip chip, where the die is flipped during packaging. This increases the performance of the device but is also more costly.

The pictures to the right display an IBM semiconductor package, with a close-up displaying wire bonding.

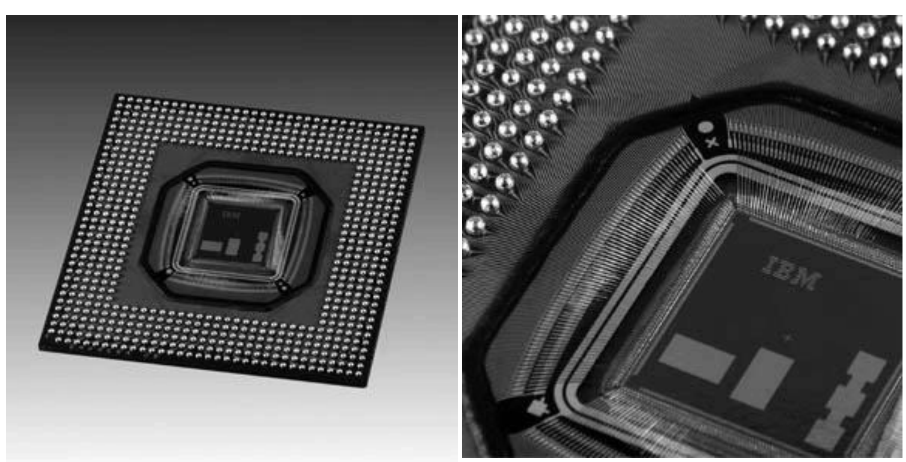

Source: IBM

Manufacturing: CMOS

CMOS (Complementary Metal Oxide Semiconductor) is the most widely used type of semiconductor design and manufacturing process. CMOS processes use standard silicon wafers and combine both positive and negative transistors.

In general, devices fabricated in CMOS will be **cheaper** to manufacture and will consume **less power** than other devices. Designs done in CMOS are also **easier to scale** to smaller transistor sizes.

Most semiconductor designers and manufacturers will use standard CMOS wherever possible. However, certain high-performance or specialized applications require the use of **bipolar** or **compound semiconductor** manufacturing processes.

CMOS (pronounced "see-mos"), which stands for complementary metal oxide semiconductor, is the most widely used type of semiconductor design process in production today. The CMOS process uses standard silicon wafers and combines PMOS (positive) and NMOS (negative) transistors in specific ways, so that the final product consumes less power than PMOS-only or NMOS-only circuits. Devices fabricated in CMOS are cheaper to manufacture and consume less power than devices fabricated using other processes, and CMOS is therefore generally used whenever possible.

Other types of manufacturing involve certain compounds more suited to specific types of applications, primarily for high-speed communications, power management, or amplification. Silicon germanium (SiGe) is a commonly used compound for high-speed physical layer devices in wireline communications. Gallium arsenide (GaAs), is used in wireless handsets and set-top boxes, as it is better suited to high-voltage RF (radio frequency) applications. Gallium arsenide, and Gallium nitride (GaN), are also used in optoelectronic devices for wireline communications. Other exotic semiconductor processes include silicon bipolar (BiCMOS), indium phosphide (InP), indium gallium phosphide (InGaP) and indium gallium arsenide (InGaAs). In general, these specialized processes are more costly to design and manufacture than standard CMOS and are more difficult to scale down to smaller process technologies (see the section on lithography starting on the next page). Therefore, wherever possible, designers tend to choose CMOS over these processes.

Although CMOS should represent the bulk of manufacturing in the near future, we note that several large IDMs, including Intel, IBM, AMD and Texas Instruments, have developed a process known as strained silicon, which allows for the use of silicon germanium transistors alongside standard CMOS transistors.

Manufacturing: Lithography

Lithography, also known as **geometry** or simply **process technology**, describes the level of **"smallness"** the manufacturing process can achieve, which determines the size of the transistors on the die.

Smaller transistors mean:

1) electrons move quicker between them → ***faster*** devices

2) more transistors per square inch → smaller die sizes → ***cheaper*** devices

3) less electricity needed to power them → ***lower power*** consumption

---

Lithography, as mentioned in the section on wafer processing, is a general term used for the set of processes that transfer the transistor array and interconnect design from the mask onto the silicon wafer. The level of "smallness" the manufacturing process can achieve is therefore determined by the lithography equipment; this in turn determines the size of the transistors on the die. The industry continually shrinks transistor size, also known as feature size, by moving to more advanced lithography equipment. On average, the industry moves to a new process geometry node every two years; however, this cadence has slowed in recent years.

The advantages of moving to smaller lithography are threefold. First and foremost, smaller transistors mean faster devices. The smaller the transistor, the less time it takes for the electrical charge to move across it. Second, smaller transistors mean smaller overall die size, which means more devices per wafer, which translates into lower cost per device. The ability to integrate additional logic, analog, discrete and memory content more cost-effectively stems from this as well. Third, smaller transistors mean lower power consumption and therefore cooler running devices.

In order to achieve smaller geometries, manufacturers must invest in new lithography equipment while chip designers must alter their designs to take advantage of the new transistors. Certain technology challenges also need to be addressed (in addition to the normal yield issues that result from any new manufacturing process), such as current leakages, increased source and drain resistance and difficulty in sending thinner laser light through glass. Most recently, the industry has required innovations such as copper interconnect and low-K dielectrics in order to meet these challenges; future innovations include high-K dielectric and extreme ultra-violet (EUV) lithography.

Manufacturing: Lithography

Moore's Law

Advancement towards finer lithography, which enables smaller and smaller transistors, is the primary driver behind the dramatic improvement in semiconductor device performance over the past few decades.

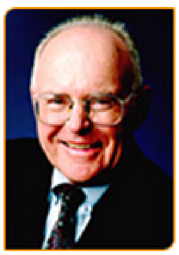

Moore's Law states that the number of transistors on a chip doubles about every two years. The phenomenon was first noticed by Intel founder Gordon Moore, and has held basically true for the past 40 years or so.

In addition to the doubling of transistor density, each new lithography generation usually brings 0.7X minimum feature scaling, 1.5X faster transistor switching speed, reduced chip power, and reduced chip cost.

The continuing push toward more advanced lithography, which enables smaller and smaller transistors, is the primary driver behind the dramatic improvement in semiconductor device performance over the past few decades.

These advancements drive Moore's Law, which states that the number of transistors on a chip doubles about every two years. Intel co-founder and industry legend Gordon Moore first noted the trend in DRAM densities and the law has held true for the past 40 years or so. Recently, however, there has been increasing debate on the relevance/cadence of Moore's law as node migration has slowed, with many industry players looking beyond transistor density for performance improvements.

Moore's Law also drives chip performance. In addition to the doubling of transistor density, each new generation usually brings 0.7X minimum feature scaling, 1.5X faster transistor switching speed, reduced chip power, and reduced cost.

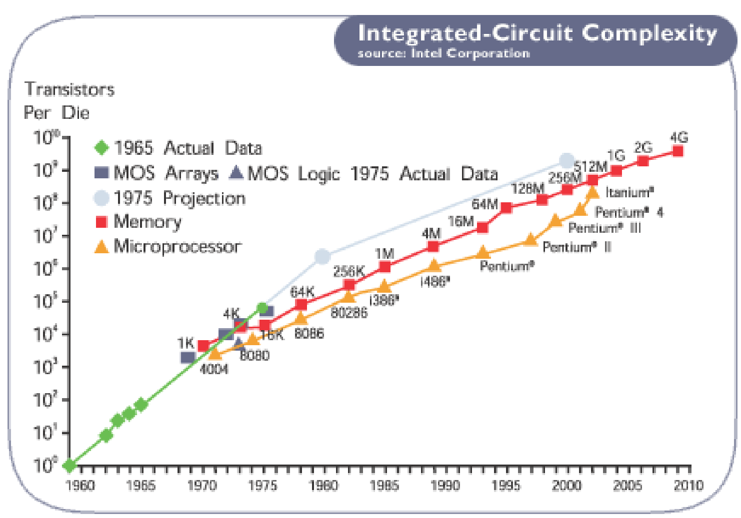

Source: Intel Corporation

Manufacturing: Strategies

Integrated Device Manufacturers (IDMs) manufacture their own devices. Typically, these are mature vendors dealing in high volume products. IDMs have lots of operating and cyclical leverage.

Examples: Intel, Texas Instruments, Samsung, Toshiba, Renesas

Fabless companies design devices themselves but contract the manufacturing to others. These companies are more focused on design, and enjoy more margin and earnings stability as well as lower capital requirements.

Examples: Broadcom, Qualcomm, NVIDIA, Xilinx, Marvell, MediaTek

Foundries and **Assembly and Testers** are specialized third-party manufacturers that perform wafer fabrication or back-end processing for others on a contract basis.

Examples (foundry): Taiwan Semiconductor, UMC, Global Foundries
Examples (assembly and test): ASE, Amkor, Siliconware Precision

---

Semiconductor companies can be divided along manufacturing lines into three segments: IDMs, fabless and foundries. IDMs design and manufacture their own devices, fabless companies design their devices but don't manufacture them and foundries perform manufacturing but no design. Increasingly, IDMs are moving toward hybrid strategies, retaining some manufacturing in-house while contracting out for some additional capacity. These hybrid strategies offer more production flexibility while minimizing the level of critical capacity investments. Others have given up on manufacturing altogether and are going fabless.

<u>Front End Manufacturing Strategies</u>

Integrated Device Manufacturers

Integrated device manufacturers, or IDMs, design and manufacture their own devices. Typically, these will be the more mature semiconductor manufacturers that deal in very high-volume products. These vendors also typically consider manufacturing expertise a key competitive advantage and a core competency for their businesses, either for technology reasons (e.g., they want to be first to market with next generation lithography, or they have specialized processes, etc.), to leverage economies of scale, to control their supply, or for a host of other reasons.

IDMs enjoy high margins in the boom part of the cycle as their fixed costs do not change, but suffer low margins in the trough of the cycle as capacity goes unused. This effect should not be underestimated, as fixed costs (including depreciation and fixed material and labor costs) can run as high as 80% of semiconductor manufacturing costs. IDMs must also maintain large capital expenditure budgets to stay ahead of the competition in order to ensure that they do not suffer performance or cost disadvantages; consequently, their ongoing cash needs are greater.

Key examples of IDMs include Intel, Samsung, Micron, as well as many Japanese semiconductor suppliers. We also consider partial-outsourcers like Texas Instruments, STMicroelectronics, Infineon, and NXP to be IDMs, as manufacturing is still a big part of what these companies do.

Fabless

Fabless companies are those that design devices themselves but contract the manufacturing to third parties. These vendors typically focus on design, leaving manufacturing competencies to others. The benefits of a fabless model include a lack of fixed costs, which translates into lower capital investment requirements, more consistent margins and better earnings predictability. Fabless companies have the opportunity to gain share vs. IDMs in the downturn (since wafer costs for them will decline while unit costs for IDMs rise with declining output) and will lose share at the peak of the cycle (because of the opposite effect). This effect can be neutralized if all direct competitors are fabless.

Key fabless vendors include Broadcom, Qualcomm, Marvell, Xilinx, NVIDIA, Mediatek, and PMC-Sierra. The list of fabless companies has been growing nicely, as the vast majority of semiconductor start-ups are fabless due to the lower capital requirements of a fabless model. Also, a number of IDMs have gone fabless in the past few years, including Agere, Avago (acquired by Broadcom), and Vitesse (acquired by MSCC).

Foundries

Foundries are specialized third-party manufacturers that perform wafer fabrication for others on a contract basis. Foundries do not design any devices themselves; rather they manufacture the designs of fabless companies and IDMs. Key foundries include Taiwan Semiconductor (TSMC) and United Microelectronics (UMC) in Taiwan, GlobalFoundries in the US, DongbuAnam in Korea, Semiconductor Manufacturing International Corp (SMIC) in China, and IBM Microelectronics in the US.

Hybrid Strategies

As mentioned above, many semiconductor companies use a hybrid manufacturing strategy. Most IDMs use external foundries for at least some of their production, usually when in-house capacity becomes tight or for specific leading-edge devices. This trend accelerated at the beginning of this decade, as IDMs found themselves burned by too much capacity in the 2001 downturn. As mentioned above, big-name IDMs like Texas Instruments, STMicroelectronics, Infineon, and NXP all outsource a significant portion of their production.

Others use a "fab-light" strategy, outsourcing most of their manufacturing but keeping some in-house. Usually, the in-house facility will be used either for older, more mature manufacturing or for specialized processes.

Another hybrid strategy, though less common, is where IDMs manufacture their own designs and also utilize the excess capacity as a foundry for additional revenue.

Back End Manufacturing Strategies

On the assembly and test side, the same three types of companies exist: those with internal facilities, those that outsource, and specialty assembly and test outsourcers. We note that very few IDMs will do 100% of the packaging themselves, as packaging types in the industry number in the hundreds and most IDMs are unwilling to invest in all the necessary equipment. Further, many IDMs (particularly high performance IC manufacturers) will focus on leading edge packaging technologies in-house, but will leverage low-cost outsourcers in Asia for the mainstream processes. Still others (particularly low cost producers) will do most of the bulk processes in-house but will contract out some of the higher performance packages.

Test functions are often conducted in the same facility as packaging and assembly. However, many fabless companies will outsource the packaging of their devices but will perform testing in-house; this allows them to more quickly identify design flaws.

The outsourced assembly and test market is less consolidated than the foundry market, with key vendors being ASE Test, Amkor Technology, Siliconware Precision, and STATS-ChipPAC. Almost all test activity is conducted in Asia, as it is more labor-intensive than front-end processes and also carries less stringent clean room requirements.

Manufacturing: Strategies

Foundries

The rising costs of semiconductor factories have made it prohibitively expensive for all but the largest vendors to do their own manufacturing. This has accelerated the trend toward outsourcing. The major foundries—TSMC, UMC, SMIC and Global Foundries—therefore nicely outpaced industry revenue growth for the past several years.

Dedicated Foundry Revenue

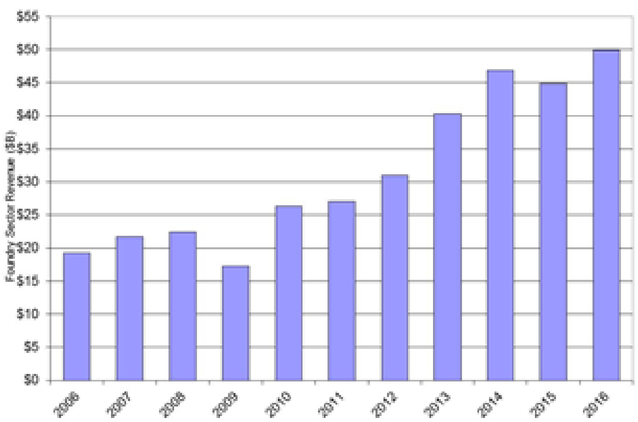

Source: Gartner, WSTS, Oppenheimer & Co.

2016 Dedicated Foundry Market Share

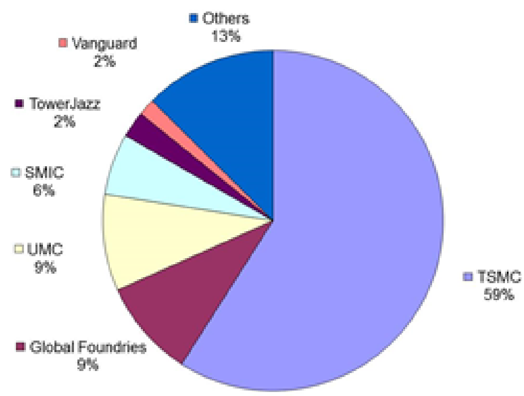

Source: Gartner Oppenheimer & Co.

The rising costs of semiconductor factories and equipment have made it prohibitively expensive for all but the largest semiconductor vendors to do their own manufacturing. This has forced many high-volume players to search for partners to do their wafer fabrication. Fabless designers have also grown in importance, and even control certain sub-segments of the semiconductor industry (graphics, PLDs, etc.); and since they are 100% outsourced, the foundries have had to grow to meet their supply needs.

At the same time, the major foundries have done an excellent job of ramping leading-edge capacity with world-class yields. The traditional top three foundries—TSMC, UMC, and GlobalFoundries—were all very aggressive in getting to 0.13-micron and have carried that through to 40-nanometer, 28nm, and 16-nm. Shanghai foundry SMIC has also gained a lot of ground with logic, memory, and fingerprint sensor producers and is considered a top-tier foundry. The incremental step up in cost to semiconductor players at lower nodes has likewise made it increasingly difficult for smaller and lower volume semiconductor suppliers to compete. TowerJazz grew 30% Y/Y in 2016 to reach the sixth highest market share due to strong RF, power management, and CMOS image sensor growth.

Revenue growth trends tell the story well: foundry revenue is up more than 2.5x since 2006. Going forward, we expect foundry revenue growth to continue to outpace the overall industry, though the rate of growth will likely decelerate.

In terms of market share, TSMC was the leading foundry provider in 2016, with 59% overall revenue share. TSMC has been a market share gainer over the past six years as a result of its advantage at the most advanced process nodes. GlobalFoundries and UMC followed with 9% each. SMIC and TowerJazz held 6% and 2%, respectively.

Manufacturing: Strategies

Assembly and Test Services

The trend toward outsourcing has also benefited back end providers, though not as dramatically as the foundries. IDMs often keep back end production in-house because they can realize lower costs, though fabless companies usually outsource. Top assembly and test houses include ASE, Amkor, Siliconware Precision, and Jiangsu

Assembly and Test Revenue

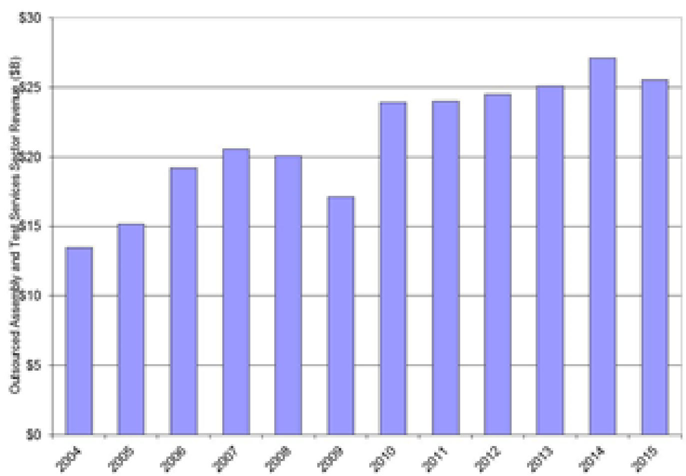

Source: Gartner, WSTS, Oppenheimer & Co.

2015 Assembly and Test Market Share

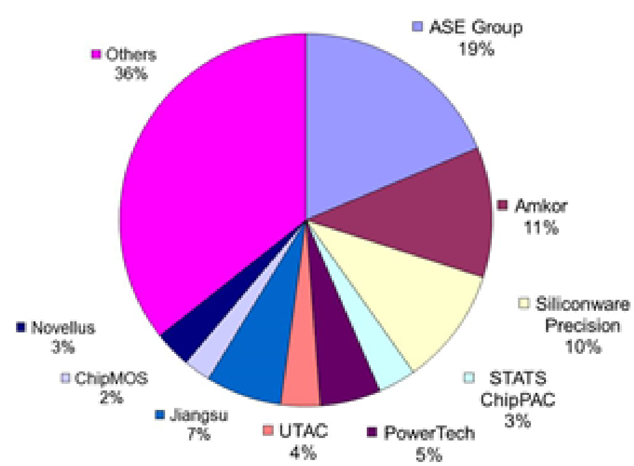

Source: Gartner, Oppenheimer & Co.

Assembly and test providers have also benefited from the trend toward outsourcing. The effect has not been quite as strong as we have seen with the foundries as some IDMs have opted to keep back-end production in-house because they can realize lower costs and the capital investment required is not nearly as large.

As mentioned above, however, very few IDMs do 100% of the packaging themselves, as packaging types in the industry number in the hundreds and most IDMs are unwilling to invest in all the necessary equipment. Many IDMs (particularly high performance IC manufacturers) focus on leading-edge packaging technologies in-house, but leverage low-cost outsourcers in Asia for the mainstream processes. Still others (particularly low-cost producers) do most of the bulk processes in-house but contract out some of the higher performance packages.

The outsourced assembly and test market is less consolidated than the foundry market. Almost all assembly and test activity is conducted in the Far East, with the top two vendors in 2015 being ASE (Taiwan) and Amkor (Korea), with 19% share and 11% share, respectively. Rounding out the top vendors were Siliconware Precision with 10% share and Jiangsu with 7% Other key vendors include PowerTech, UTAC, STATS ChipPac, ChipMOS, and Novellus.

Manufacturing: Geographic Centers

The largest geographic centers for semiconductor front end production are Japan and Asia, each about 1/3 of total production, with the final 1/3 split roughly evenly between North America and Europe.

Asia has grown in importance over the past two decades: Korea is a major center for memory, displays, and commodities, and Taiwan hosts the largest foundries as well as some memory producers. China is also ramping up with several fabs under construction.

	Japan	North America	Europe and Israel	Taiwan	Korea	China
Key IDMs	Renesas Toshiba Sharp Cypress Panasonic Sony Micron	Intel Texas Instruments NXP IBM Micron Analog Devices	Infineon STMicro NXP AMD Intel Bosch	Winbond Nanya Powerchip ProMOS	Samsung Hynix Magnachip	
Major Foundries	Toshiba Seiko-Epson Fujitsu	IBM TowerJazz GlobalFoundries	X-Fab TowerJazz	TSMC UMC	Magnachip DongbuAnam	SMIC ASMC Grace
Key Products	Microcontrollers Discretes Memory Logic, ASSPs Analog	MPUs, MCUs, DSPs Logic, ASSPs Analog Memory Discretes	Logic, ASSPs Analog MPUs Discretes Memory	Logic, ASSPs Memory Display Drivers	Memory Logic, ASSPs Display Drivers	Memory Logic

The largest geographic centers for semiconductor front-end production are Japan and Asia. Japanese electronics OEMs have historically been vertically integrated and thus retain fabs for their own semiconductor production. They are also large suppliers of microcontrollers, discretes, general purpose logic, and analog ICs to the broader market. DRAM, a big priority in Japan in the 1980s, has shifted away from Japan for the most part, but flash has ramped up for both NOR and NAND.

Asia ex-Japan has clearly grown in importance over the last two decades, representing ~70% of semiconductor revenue. Korea is a major center for memory, displays and commodities, and electronics giant Samsung continues to invest and grow in importance. Taiwan (12% of total) hosts the largest foundries as well as some memory producers. China is also ramping up, mostly with foundry production for memory and logic.

The US (Americas equal 14% of total) is still the most important region for high-end logic, with most CPU and DSP production as well as microcontrollers and ASSPs. Custom ASICs, a wide array of memories, analog ICs and commodities also have a large presence. Europe (10% of total, including fabs there owned by U.S. companies) has a similar mix; Infineon, STMicroelectronics, and NXP are large diversified players that compete across these categories, and Intel runs large microprocessor fabs in the region (Ireland and Israel).

Note that back-end factories are more weighted toward Asia due to the large labor component. Back end facilities are spread throughout the region, not only in Taiwan and Korea but also in China, Malaysia, Singapore, and elsewhere.

Semiconductor Equipment

The semiconductor industry has a large and relatively independent supporting industry that supplies equipment for front and back end processing.

Top equipment suppliers include Applied Materials, ASM Lithography, Tokyo Electron, Lam Research, KLA-Tencor and Dainippon Screen.

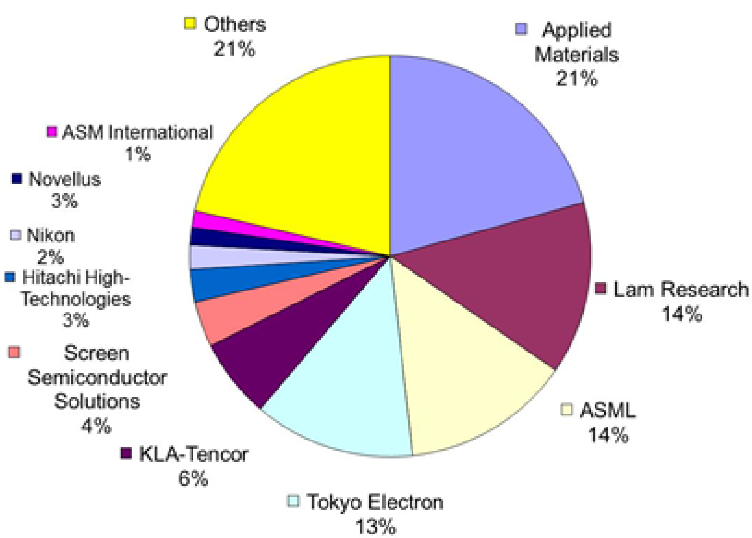

Source: Gartner, Oppenheimer & Co.

Front End	Process Control
Applied Materials ASML Tokyo Electron Lam Research KLA-Tencor Nikon	KLA-Tencor Hitachi Applied Materials Nanometrics Hermes
Die-Level Manufacturing	**Automated Test**
ASM Intl. Disco BE Semi Shinkawa	Advantest Teradyne Tesec

The semiconductor industry has a large and relatively independent supporting industry that supplies wafer fabrication, assembly, and test equipment as well as factory automation and process control software and hardware. Few IDMs or foundries make their own equipment, though some of the larger manufacturers will partner with equipment suppliers to help define the functionality of next-generation tools.

In 2016, the semiconductor capital equipment industry expanded to \$37.0 billion in sales. Approximately 84% of this total was for front-end fab equipment, which includes CVD and sputtering deposition, diffusion/oxidation and RTP, lithography, ion implantation and process control. Another 11% was for process control, while the remaining 5% of equipment spending was for back-end equipment, including packaging, assembly and automated test equipment. The exhibit to the right displays this breakdown.

Semiconductor Equipment Spending, 2016

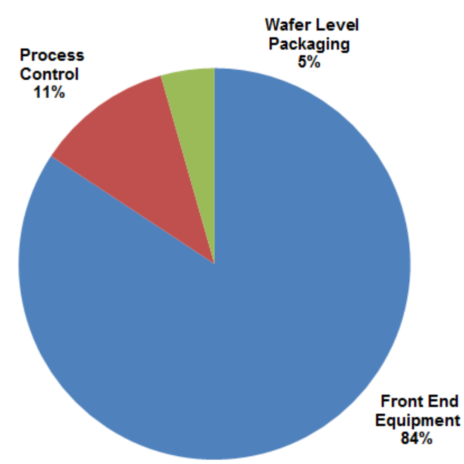

Source: Dataquest, Oppenheimer & Co.

Top front-end equipment suppliers include Applied Materials, strong in deposition, RTP (rapid thermal processing), CMP (chemical mechanical planarization), and etch; Tokyo Electron, prominent in thermal diffusion, photoresist processing, and etch; ASML Lithography, dominant in lithography scanners; and Lam Research in etch. Together, the top four front-end vendors account for 63% of total equipment.

On the process control front and automation front, top vendors include KLA-Tencor, strong in wafer and reticle inspection/metrology; Hitachi, also strong in wafer metrology; Applied Materials mostly in defect review.

The exhibits below display fab equipment market shares for front-end and back-end equipment as well as automation and process control.

2016 Front End Equipment Market Share

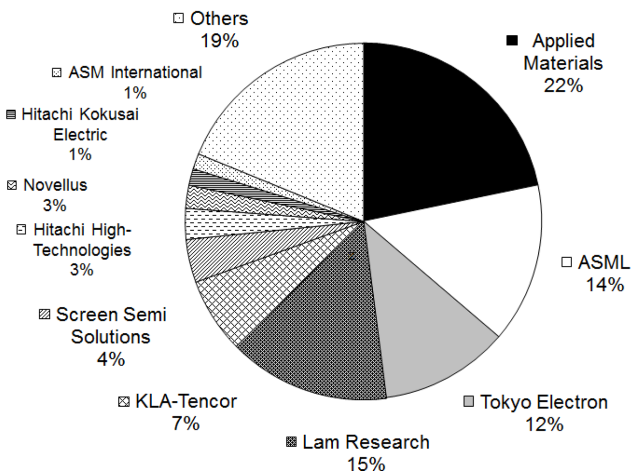

2016 Process Control Market Share

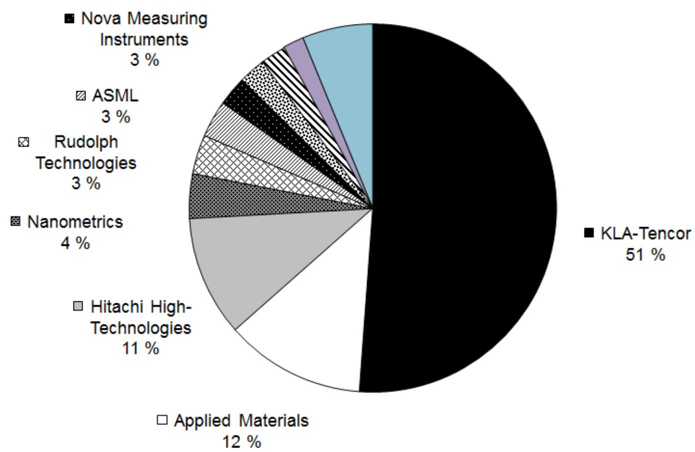

Source: Gartner, Dataquest, Oppenheimer & Co.

Section II: The Semiconductor Market

This section discusses the semiconductor market, industry fundamentals, and the semiconductor cycle

Topics include:

Industry Basics

The Semiconductor Cycle

Fundamentals

Industry Basics: 40 Years of Growth

As the enablers of electronic hardware, semiconductors have been at the heart of the technology revolution. The market has grown significantly over the past 40 years, exceeding $300 billion in revenue for the first time in 2013.

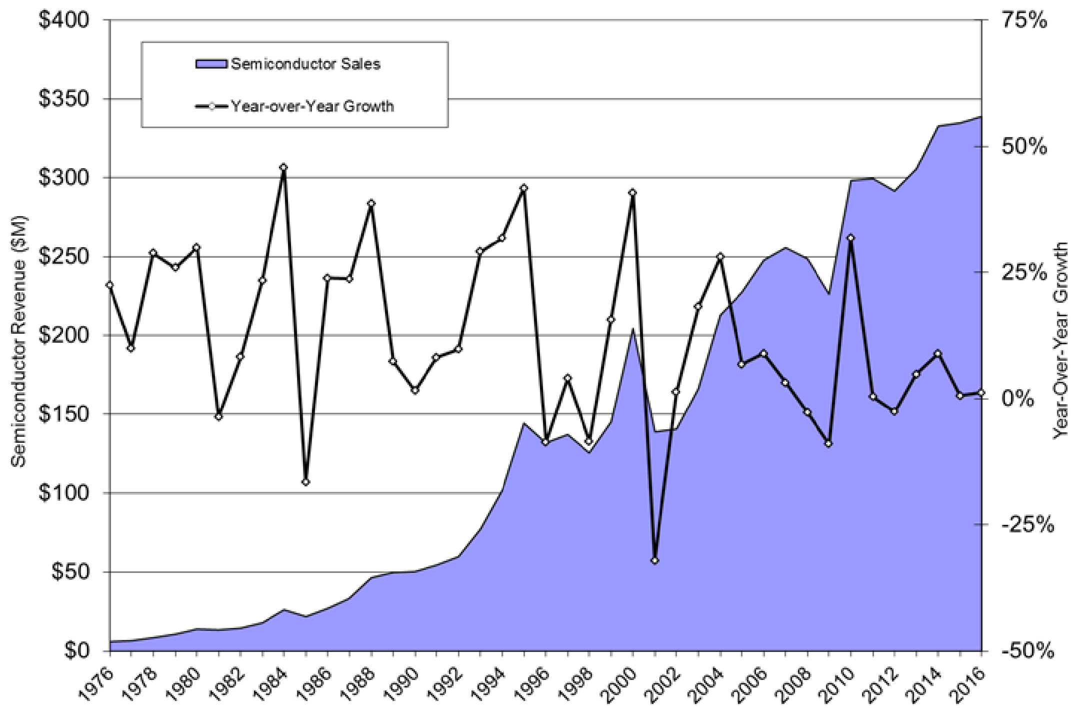

Source: WSTS, Oppenheimer & Co.

As the enablers of electronic hardware, semiconductors have been at the heart of the technology sector. The market has grown tremendously in the past 40 years; however, growth has decelerated over the past decade as some segments have become commoditized, and pricing has come under pressure from low-cost Asian suppliers.

As mentioned, growth has been slowing, as the industry has grown and matured. We have plotted a 5-year moving average of the industry's 5-year CAGR on the graph to the right. After hovering in the mid-to-high teens for more than two decades, we see the average dropping in the post-bubble and financial crisis periods. We would expect the CAGR to remain in the single digits, growing at a GDP+ percentage for the forseeable.

5-Year Moving Average of 5-Year CAGR

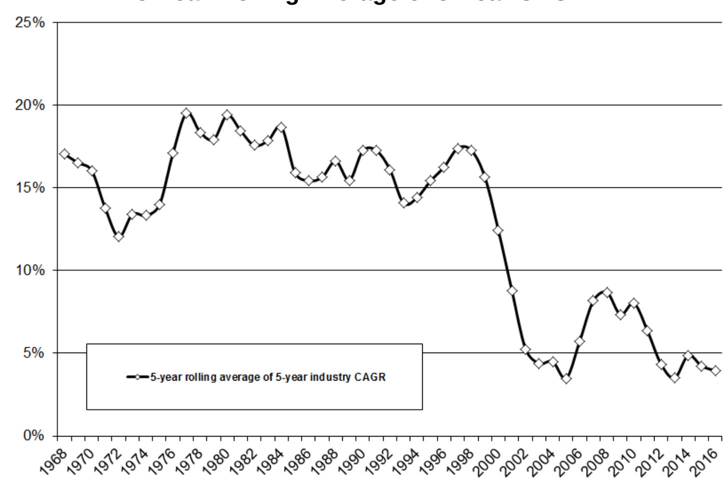

Source: WSTS, Oppenheimer & Co.

Industry Basics: Demand

Semiconductor industry demand is closely tied to the demand for electronics. The largest consumers of semiconductors are the computing/data processing, communications, and consumer electronics markets. Semiconductor revenue has also been growing as a percentage of electronics revenue as chipmakers have absorbed more and more system IP into their devices.

Electronics Revenue By End Market

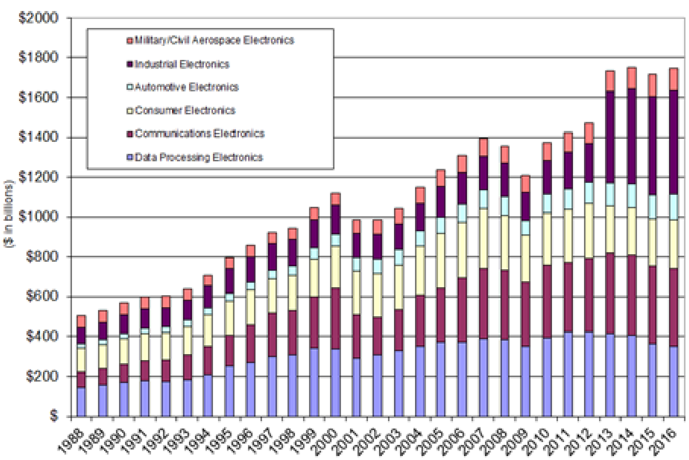

Source: Gartner, Oppenheimer & Co.

Semi Revenue as a Percentage of Electronics

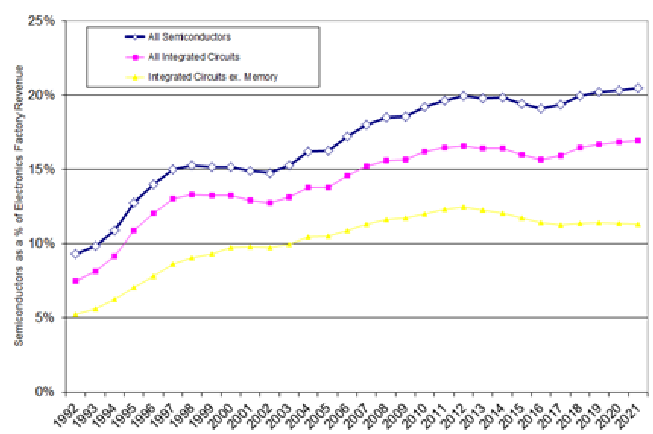

Note: Data displayed is five-year average to smooth price volatility.
Source: Gartner, WSTS, Oppenheimer & Co.

Semiconductor companies support the large and diverse electronics industry. Semiconductor demand is thus a function of the demand for electronic equipment along with the percentage of electronic equipment bill-of-materials devoted to semiconductors.

Electronics Demand

The electronics industry totaled about $1.7 trillion in 2016, up 2% from 2015. By end market, the most important applications are computing/data processing, communications and consumer electronics. In 2016, 29% of electronics revenue was derived from Industrial electronics, while 20% was data processing, which includes PCs and servers, PC peripherals, tablets, displays and storage. Another 22% was derived from communications, including smartphones, networking equipment, telecom and wireless infrastructure, and voice and data access equipment. Fourteen percent was derived from consumer electronics, including digital TVs, audio/visual equipment, game consoles, set-top boxes, appliances and white goods. Another 8% was derived from automotive, and the remaining 7% from military and aerospace.

The exhibits on the following page detail electronics revenue by application.

2016 Data Processing Equipment Revenue (20%)

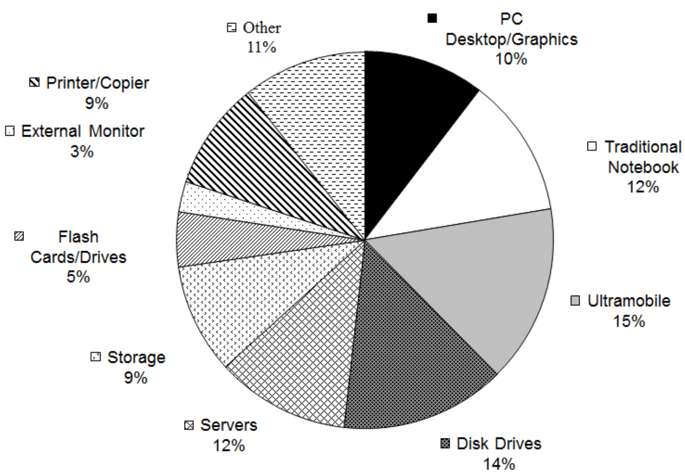

2016 Communications Equipment Revenue (22%)

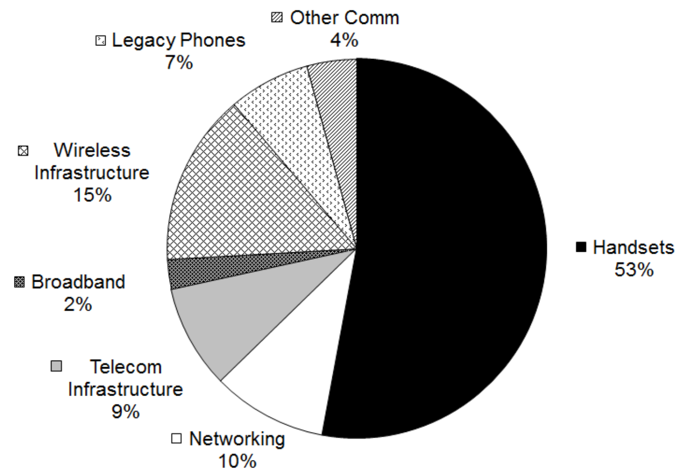

2016 Consumer Electronics Revenue (14%)

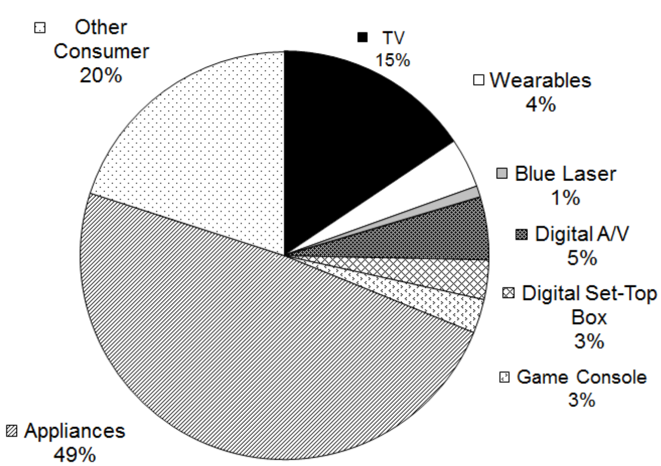

2016 Other Electronic Equipment Revenue (44%)

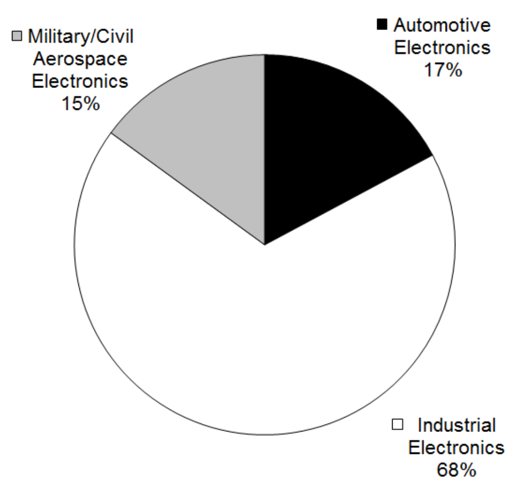

Source: Gartner, Dataquest, Oppenheimer & Co.

Semiconductors as a Percentage of Electronics

On top of electronics demand, semiconductor revenue has been growing as a percentage of the bill-of-materials and of total electronics revenue. In 2014, this ratio was approximately 19%, having climbed steadily from the 10% range in the early 1990s. This increase has been due to a number of factors, including:

- Chipmakers have absorbed more and more system IP into their devices. Chip prices increasingly include software components and other costs.
- Outside of the memory market, chipmakers have for the most part retained pricing power in their devices, whereas OEMs have been subject to tough competition, pressuring their pricing.
- OEMs have been able to take costs out of the box through outsourcing, just-in-time inventory, etc. Chipmakers have benefited from Moore's Law, but fab costs have continued to rise, as have R&D costs.
- OEMs have increasingly diversified their revenue streams by charging their customers for software, services, maintenance contracts and other non-equipment products. This has allowed them to charge lower margins on equipment. Semiconductor vendors, in contrast, continue to charge only for actual physical product (though they provide some software, this is usually embedded in the price of the IC).

Industry Basics: Seasonality

The semiconductor industry is somewhat seasonal, based on demand patterns for electronics. Typically, 47%-48% of sales fall in the first half of the year, while 52%-53% fall in the second half. On a quarter-to-quarter basis, sales typically fall in Q1, rise a bit in Q2 and rise more strongly in Q3 and Q4. In general, devices sold into consumer channels (CE, handsets, PCs) are typically more seasonal than those sold to commercial customers.

Percentage Of Yearly Sales By Quarter

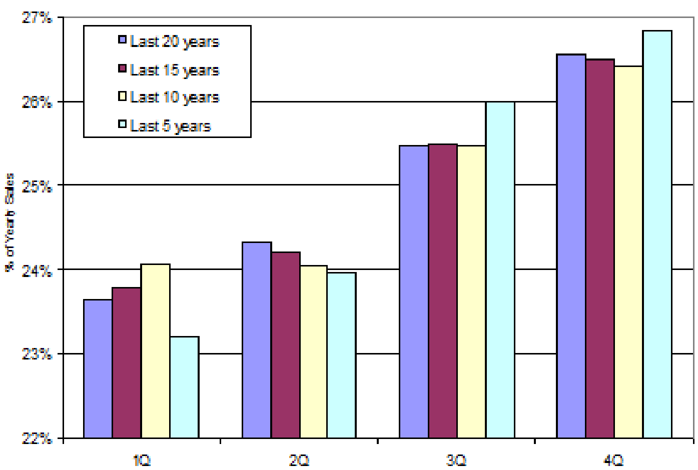

Source: WSTS, Oppenheimer & Co.

Historical Quarterly Sequential Growth

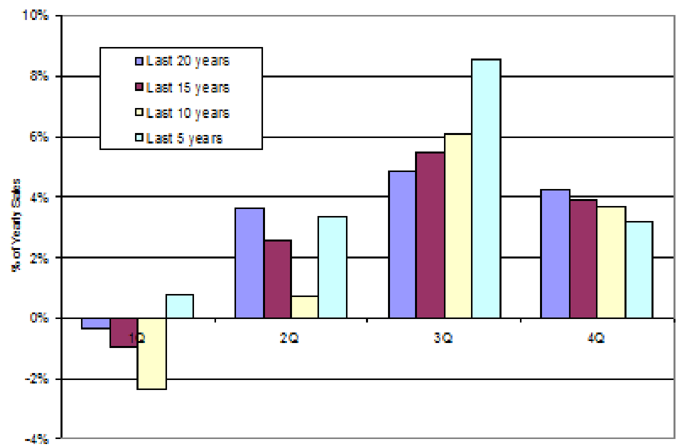

Source: WSTS, Oppenheimer & Co.

The semiconductor industry is somewhat seasonal, based on demand patterns for electronics but with a slight lead due to the time it takes to build, test and ship the end product (semiconductors will usually go on the board in the early stages of production). These demand patterns tend to favor Q4, which sees a bump in demand for the holiday selling season; and also for Q3, which sees the “holiday build” as well as a somewhat less significant “back to school” season. This is especially true for devices sold into the consumer channel, including consumer electronics, handsets and PCs. Demand for semiconductors sold into enterprise, service provider, and industrial applications will see smoother demand, though there can sometimes be a budget flush at the end of the year.

Typically, 47%-48% of sales fall in the first half of the year, while 52%-53% fall in the second half. On a quarter-to-quarter basis, sales typically fall in Q1, are up modestly in Q2, and rise more strongly in Q3 and Q4. Seasonality, however, is changing as the consumer electronics industry gravitates toward mobility. We have witnessed a dampening of the Q4 and Q1 sequential changes as electronics consumption in Asia becomes more meaningful. Asian consumers have a holiday selling season around Chinese New Year, which usually occurs in late January or early February.

The Traditional Supply-Based Cycle

In the past, the semiconductor industry has historically been characterized by supply-based cycles lasting 4-6 years on average. Wafer capacity was the key variable driving these cycles. Demand fluctuations were of course a factor, but not nearly as important as the state of supply.

	Phase One "The Downtum"	Phase Two "The Recovery"	Phase Three "The Expansion"	Phase Four "The Peak"
Fundamentals				
Macroeconomics				
Economic activity	Recession or Slower Growth	Picking up	Expansion	Stable growth
Pricing				
Logic	Decreasing	Decreasing to stable	Increasing	Stable
Analog	Decreasing	Decreasing	Stable to increasing	Increasing
Commodities	Decreasing rapidly	Decreasing	Increasing	Peaking
Lead times				
Analog/logic	Contracting	Stable	Stable to expanding	Expanding to peaking
Commodities	Contracting rapidly	Contracting to stable	Expanding	Peaking - on allocation
Capital additions				
Production line additions	No (over-capacity)	Some	Yes	Yes (too many)
New capacity additions	No (over-capacity)	None	Yes	Yes (too many)
OEM/distributors				
Orders	Falling rapidly	Picking up	Strong	Frenzied
Double ordering	No	No	No	Yes
Order cancellations	Accelerating	Falling	Very few, if any	Limited
Inventory levels	Too high	Low	Too low	Building
Vertical end markets				
Sell-through	Even slower	Growing	Strong	Slowing
Forward outlook	Cautionary	Cautionary	Cautiously optimistic	Optimistic

Source: Oppenheimer & Co.

The semiconductor industry has historically been characterized by supply-based cycles lasting four to six years, on average. Wafer capacity was the key variable driving these cycles. Demand fluctuations were of course a factor, but not nearly as important as the state of supply. This traditional cycle can generally be broken down into four key phases:

- **Downturn**: This phase follows the peak of the previous cycle and is usually sparked by slowing economic conditions coupled with oversupply. Orders dry up and cancellations appear. OEMs start to work through too-high inventories. ASPs begin to fall as capacity built during the boom period pressures pricing against weaker demand. IDMs and foundries cancel pilot line additions and even close down or consolidate trailing-edge fabs.
- **Recovery:** The recovery is a period of stabilization following the free-fall of the downturn. Utilization rates level off after the drop, and manufacturing rationalization continues. Prices stabilize but remain at low levels. Orders begin to pick up and cancellations abate. Inventories subsist at low levels after a long burn-off.
- **Expansion:** Eventually, supply and inventory correct sufficiently to drive a new expansion. Economic activity drives end demand, which causes stronger orders. Utilization and then prices begin to rise and IDMs begin to cautiously add capacity to meet increasing output forecasts. Lead times stretch and order visibility improves.
- **Peak:** Eventually, enough capacity is added to meet surging demand. Some of this demand is based on double ordering, and so as vendors catch up with deliveries, inventory builds in the channel. Prices and lead times start to peak out. Eventually, double-ordering abates as vendors realize they have too much capacity.

The Evolving Cycle

The semiconductor cycle has evolved since the tumultuous 2000-2001 boom-bust:

- The growing use of foundries by both fabless and IDMs has removed leverage from semiconductor business models
- Diversification of the semiconductor industry's end market lessens the dependence on PC cycles
- Fractured supply chain spreads risk
- Law of large numbers and slowing growth
- Industry consolidation leads to supply chain rationalization

Key Implications:

- New cycles are based more on inventory than supply
- Shorter cycles with shallower ups and downs
- Gross margin less relevant for all but a handful of producers
- Memory cycles much more independent from the broader semi cycle
- Industry revenues grow even in weak fundamental periods so long as macro economy is healthy

---

The cyclical nature of the semiconductor industry has driven vendors to restructure and re-work their business models from cycle-to-cycle. This was especially true in the period following the 2000-2001 boom-bust, which featured the traditional supply-based peak and downturn exacerbated by the overhang from excess Y2K and telecom/Internet infrastructure spending.

A key component of this evolution has been the growing use of foundries by not only fabless semiconductor vendors, but also by IDMs. Large producers such as Texas Instruments, STMicroelectronics, NXP, and Infineon began to outsource a significant portion of their production to foundries. This shifted a considerable amount of supply risk for the industry into the hands of just a few leading-edge foundries, which have proven that they can invest more prudently than individual players (who often overbuild to support projected market share gains, which often don't materialize).

As a result, the industry has seen much less over-building during periods of strong demand, and exhibited much less under-investment during weak periods. IDMs have been able to keep expensive digital logic fabs full most of the time, leveraging foundries for overages.

The net effect of the increased outsourcing has been a much more muted connection between wafer supply and other industry metrics, such as lead times, order rates, pricing, etc. IDMs view supply much in the same way fabless vendors do: as a steady component of their business rather than a sizeable capital drain which needs to be constantly adjusted up and down.

Consequently, semiconductor industry cycles are no longer primarily driven by wafer supply directly. Instead, they have become driven by an indirect aspect of supply: specifically inventories. Downturns which used to be driven by over-supply are now driven by excess inventories in the supply chain, which usually result from a slowdown in demand. Expansions are driven by inventory re-builds, which are usually a result of an increase in demand against a backdrop of a tight supply chain.

The new inventory cycle is now shorter—roughly 1-3 years from peak-to-peak—and has shallower ups and downs. The phases of the cycle are harder to identify. Industry revenues can grow in both strong (expansion/re-build) and weak (downturn/correction) periods, as long as the broader macro economy and overall demand for electronics are healthy.

We believe the primary source of the cycle's evolution lies in the wafer outsourcing detailed above, but we also see secondary drivers as well. Greater diversification of the industry's end demand has certainly contributed some additional stability. The industry used to live and die by PC cycles, since such a significant portion of semiconductor demand was driven by PCs. Today's semiconductor industry is well-balanced between computing, communications, industrial, and consumer electronics applications, while the end use is spread among enterprises, service providers, and consumers. Geographically, emerging markets have grown to represent a sizeable consumer of semiconductors, which used to be an industry dominated by the US, Japan, and Europe. This diversification leads to much smaller swings in demand based on factors in any one application, economic sector, end market, or region, and hence greater stability and less cyclicality overall.

An additional factor is the fractured nature of the modern electronics supply chain. Semiconductor vendors, foundries, OEMs, ODMs, EMS players, and distributors are each important constituents within the supply chain. Each has matured and has learned to play a stabilizing role within the greater whole. This spreads risk among a greater number of players and lessens the shock potential from any one segment of the supply chain.

Finally, the law of large numbers and a general maturity are clearly setting in for the semiconductor industry. Industry sales are now above $300B annually and global per capita consumption of semiconductors now tops $40 for every man, woman, and child (much higher in developed nations). Mid-to-high teens percentage growth rates are simply hard to come by off this high base, and the lower growth rates are also producing smaller changes year-to-year in terms of growth. This has muted a major input for the cycle: semiconductor demand.

Note that the evolution of the cycle described here is inclusive of most analog and digital ICs, with one notable exception: memory. The memory sector remains mired in the traditional supply-driven cycle, owing to 1) the vertically integrated business models that remain the norm among top players, 2) the lack of diversification in end application that continues for each memory type (e.g. DRAM still dominated by PCs, NAND by consumer), and 3) the commodity nature of most memory types, which naturally breed the supply-driven economics. It is even possible for one type of memory to be experiencing an expansion while another type is caught in a downturn, based on the supply-demand dynamics of the individual memory type.

The New Inventory-Driven Cycle

Structural changes in the semiconductor industry have produced a new cycle, one that is inventory-driven. Inventory is the key variable along with demand; wafer supply is relegated to a secondary input. Cycles are shorter (1-3 years) with shallower peaks and valleys.

	Phase One Inventory Correction	Phase Two Steady State	Phase Three Inventory Build
Fundamentals			
Macroeconomics			
Economic activity	Slower growth or recession	Stable growth	Stable growth or expansion
Sell-through in verticals	Slow	Growing	Strong
Forward outlook	Cautionary	Cautiously optimistic	Cautiously optimistic
Inventories			
On hand	Too much	Low but managemeable	Too little
Posture	Reduction	Conservative restocking	Double ordering
Pricing			
Logic	Decreasing	Decreasing to stable	Increasing
Analog	Decreasing	Decreasing	Stable to increasing
Commodities	Decreasing rapidly	Decreasing	Increasing
Lead times			
Analog/logic	Contracting	Stable	Stable to expanding
Commodities	Contracting rapidly	Contracting to stable	Expanding
Capacity			
Availability	Widely available	Widely available	Tightening, espcially at the high-end
Foundry pricing	Declining	Stable	Firm
Orders			
Order rates	Falling rapidly	Strong	Frenzied
Double ordering	No	No	Yes
Order cancellations	Accelerating	Very few, if any	Limited

Source: Oppenheimer & Co.

Structural changes in the semiconductor industry have produced a new cycle, one that is inventory driven. Inventory is the key variable along with demand; wafer supply is relegated to a secondary input. Cycles are shorter (two to three years) with shallower peaks and valleys. This new cycle can be broken down into three key phases:

- **Inventory Correction**: This phase follows the peak of the previous cycle and is usually sparked by slowing economic conditions against a backdrop of an inventory glut in the supply chain. OEMs and distributors seek to reduce inventories and reduce orders, or even cancel orders at semiconductor vendors. ASPs for more commoditized products begin to fall. The foundries (and to a lesser extent the IDMs) scale back on capacity additions until the true level of demand is determined.
- **Steady State:** The steady state phase is a period of stabilization following the steep declines of the inventory correction. Orders begin to pick up and cancellations abate. Growth rates appear more "normal," approximating growth in end equipment. Inventories subsist at low levels after a long burn-off, and may even be managed lower. Utilization rates level off after the drop in production, but lead times remain short.
- **Inventory Build:** Eventually, inventories at OEMs and distributors get too low, and they begin to re-order. Lead times begin to stretch out. A pickup in economic activity drives end demand, and in some cases causes double-ordering. Utilization improves and foundries begin to slow price declines to their customers. Prices for commodities begin to rise. Eventually, inventory builds up on excess of demand, sparking the next correction.

Fundamentals: Sales

Semiconductor industry sales reached an all time high of $339 billion in 2016. The industry had experienced four consecutive years of record sales until a contraction in 2008 and 2009 led to revenues falling below $250 billion. Structural changes during that recovery have replaced the supply-driven cycle with the inventory-driven cycle, and the industry expanded in 2010 and 2011 before contracting in 2012, primarily due to weakness in the PC market. After reaccelerating in 2013 and 2014, the industry saw muted growth in 2015 and 2016

Yearly Sales

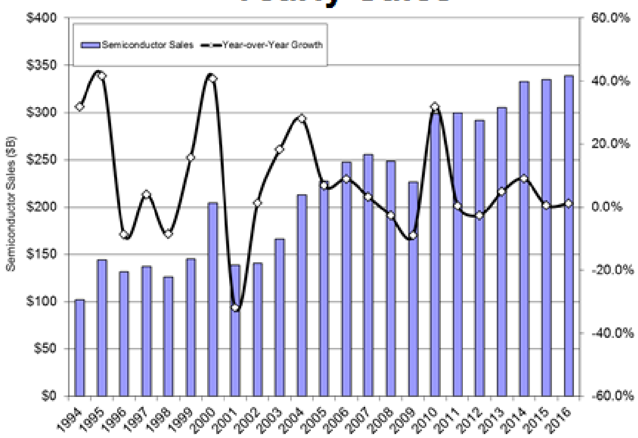

Source: WSTS, Oppenheimer & Co.

Monthly Sales

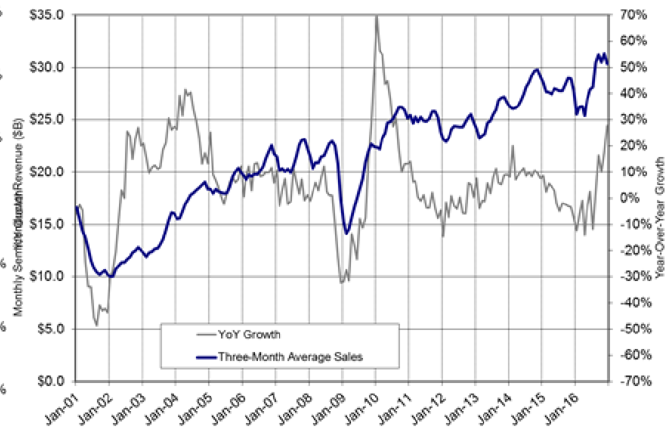

Source: WSTS, Oppenheimer & Co.

The breakout growth of electronic hardware has driven dramatic growth in semiconductor sales over the long term. Semiconductor industry sales in 2016 totaled $339 billion. After four straight years of record sales for the industry, following the long downturn and recovery period of 2001-03, the industry experienced a contraction in 2008 and 2009 given macroeconomic concerns. Semiconductor sales bounced back sharply in 2010, as end demand recovered faster than anticipated and as OEMs restocked inventory. Industry sales flat-lined in 2011 and 2012 before reaccelerating in 2013 and 2014. Subseqently, however, the semiconductor industry saw muted growth in 2015 and 2016, but appears on pace for a strong 2017 led by strong memory pricing.

The cyclical nature of the semiconductor industry has caused powerful boom periods and destructive downturns. Fundamental industry characteristics such as long production cycles, the quick pace of innovation (funded by increasingly tech-centric capital markets), volatile pricing, inventory effects and capacity builds have caused wild sales volatility year-to-year.

The 2001 downturn was so destructive that it motivated the big industry players to make real structural changes to their business models. This has caused the profile of the semiconductor cycle to change as well, as detailed on prior pages. The “new” cycle is primarily inventory-driven as opposed to supply-driven, and though the ups and downs of the cycle remain, the magnitude and duration of these swings is much reduced. As a result, we have seen year-over-year growth rates, in general, settle into a more consistent and manageable range.

Fundamentals: Units and ASPs

Unit shipments are slightly less volatile than sales and have historically had a more upward slope. Still, unit shipments are subject to double ordering and inventory corrections, and so still swing around during the cycle.

ASPs have remained volatile and generally add to the boom during periods of strong unit demand (and when inventories build) and exacerbate the downturn when demand softens (and inventories decline).

Unit Shipments

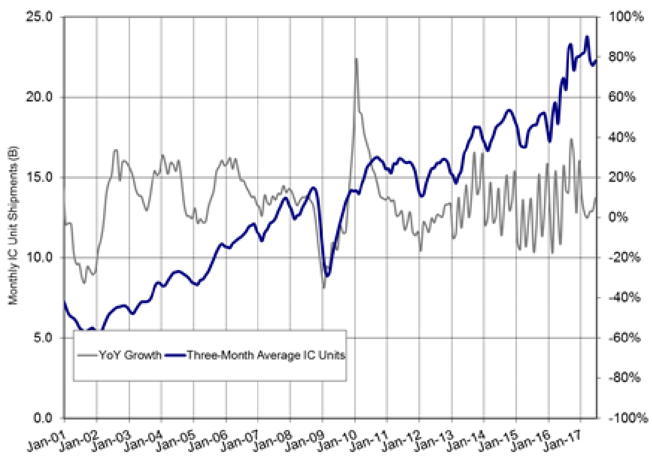

Source: WSTS, Oppenheimer & Co.

Average Selling Price

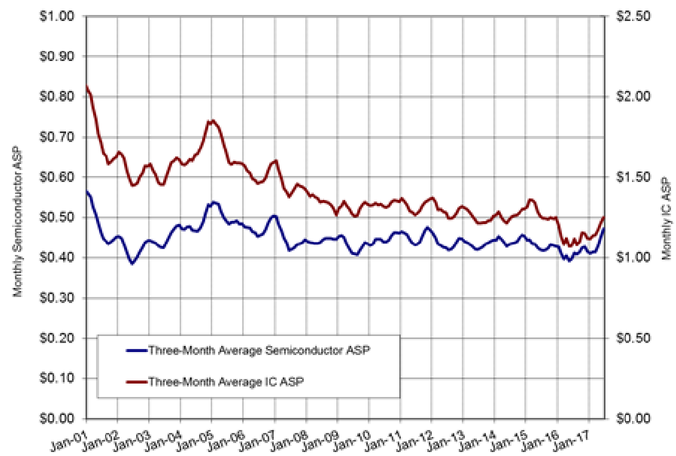

Source: WSTS, Oppenheimer & Co.

Semiconductor unit shipments have also proven volatile, though they generally show a bit more stability and a more upward long-term slope than revenues. Still, shipments are affected by periods of double-ordering and subsequent inventory corrections, which exaggerate more muted ups and downs in true end demand. As with sales, the swings in unit shipments have softened as the cycle has evolved from a supply-based cycle to an inventory–based one.

Pricing also moves around with the cycle, and in general has had a downward slope as vendors pass economies of scale and the benefits of Moore's Law onto customers. The slope has been leveling off to some degree, however, as a major long-term trend of integration has allowed semiconductor vendors to slow the pace of ASP declines on individual components.

Pricing was a major factor in supply-based cycles, adding to the boom during periods of strong unit demand and exacerbating the downturn when demand decelerates. The commoditization of memory products like DRAM and flash has added to this phenomenon as it represents such a large percentage of industry sales. Note that the evolution of the cycle has mitigated ASP volatility to some extent, as has the general maturity of the industry. Still, commodity and even logic and analog pricing can have wild swings on demand and supply fluctuations.

Fundamentals: Gross Margin

Gross margins for the industry are generally in the high-40% to low-50% range, though they compress during periods of low utilization or difficult pricing. Outsourcing trends have resulted in less margin volatility since the 2000 bubble.

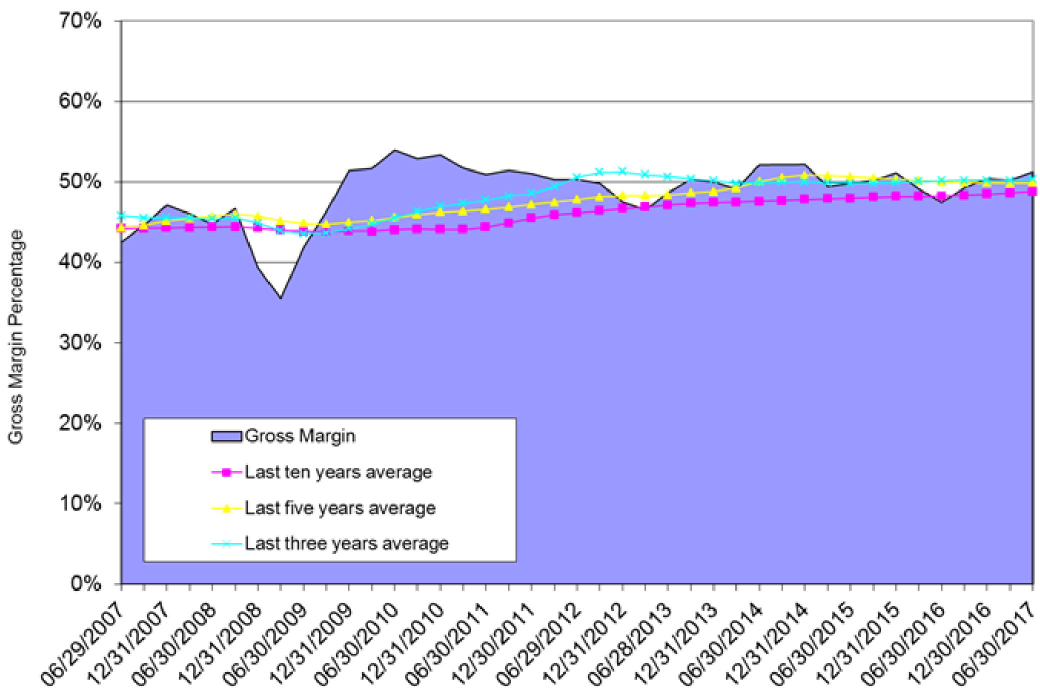

Source: FactSet, Oppenheimer & Co.

Gross margins for semiconductor companies have generally hovered in the 40-50% range, recently setlling in the high-end of the range. A hefty portion of COGS is driven by depreciation of plant and equipment, with the remainder in fixed and variable labor and material costs. In general, companies whose products are more differentiated and proprietary tend to have higher margins; this would include vendors of analog and mixed signal ICs, PC microprocessors, DSPs, FPGAs, and special purpose logic (ASSPs). Companies whose products are more commoditized and thus differentiate themselves based on their manufacturing competencies tend to have lower margins; this would include vendors of memory, microcontrollers, discretes, optoelectronics and general purpose logic.

Gross margins tend to compress rapidly during periods of difficult pricing or low utilization and expand quickly when prices firm and utilization is on the rise. The effect can be exaggerated by industry accounting practices employed by many IDMs, whereby fixed costs are allocated to inventory each quarter at an estimated standard utilization rate and excess fixed costs (positive or negative) are charged to the cost of goods line. In periods of low utilization, actual production trails the estimate, hurting gross margin in the current period. In periods of high utilization, actual production leads the estimate, leading to negative charges to cost of goods and a boost to gross margin.

Outsourcing trends have mitigated margin volatility since the 2001 downturn, and now only a handful of semiconductor vendors (mostly memory vendors) do not outsource at least 10% of their supply needs. This eliminates much of the margin leverage we are used to seeing in IDM business models. Even amid challenging conditions in the 2004 and 2006 inventory downturns, industry gross margins held above 40%. During the 2008-09 correction, industry gross margins dipped slightly below 40%. Gross margins rebounded into the 50's (%) during 2010 before normalizing backwards into the high-40% range in 2013 and the last several years, as semiconductors have become an increasingly important part and point of differentiation in high-value end-markets, including automotive and industrial.

Fundamentals: Operating Margins

Operating margins had historically averaged under 15%. Margins turned negative in 2009, and recovered above historical averages as sales rebounded more quickly than spending. Operating margins have come under increased scrutiny in recent years, as investors look for earnings leverage as overall growth stagnates. Even in a choppy, inconsistent demand environment, companies remain heavily invested in R&D, thus much of the cost savings has come on the SG&A line. We believe industry consolidation has, in part, been driven by an increased focus on operating margins as growth has slowed.

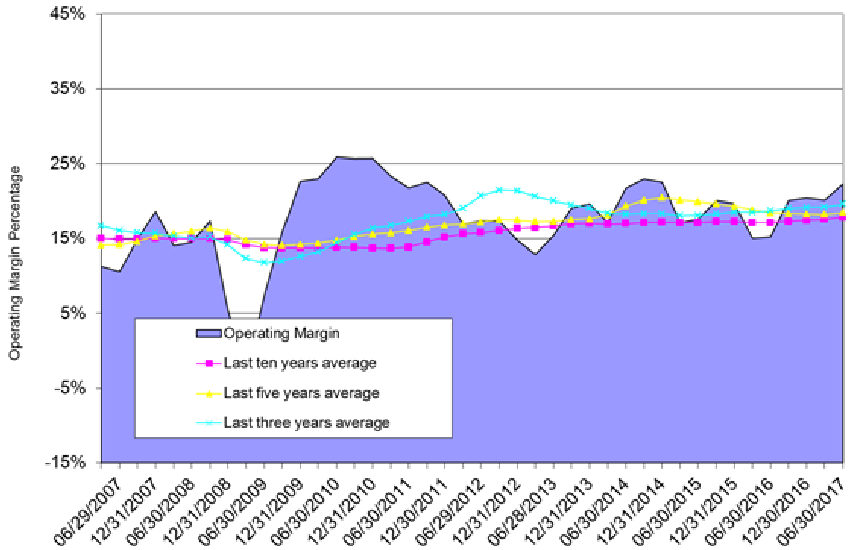

Source: FactSet, Oppenheimer & Co.

Operating margins for semiconductor companies have historically averaged in the 10-15% range, with the average company spending about 25-30% of sales on R&D and SG&A expenses. Operating margins above 20% are not uncommon, however, especially for best-in-breed companies; or even for laggards in peak periods of the cycle.

Operating margins for the industry turned negative during the 2001 downturn and again during late 2008/early 2009. Operating margins did not exceed 20% for any length of time at the peak of the earlier cycle, but rebounded above 25% in 2010. Along with declining industry sales and moderating gross margins, operating margins have consolidated toward normalized levels into 2013.

Operating margins have come under increased scrutiny in recent years, as investors push for earnings leverage as overall growth stagnates. Even in a choppy, inconsistent demand environment, companies remain heavily invested in R&D; thus, much of the cost savings has come on the SG&A line. We believe industry consolidation has, in part, been driven by an increased focus on operating margins as growth has slowed.

Many semiconductor vendors have exhibited an increased focus on maintaining profitability through a move to more variable costs. Most have undertaken some level of restructuring, including the move to more flexible outsourcing models, headcount reductions, and strategy shifts. Still, the overwhelming majority of companies remain determined to succeed by staying heavily invested in R&D and SG&A, producing sub-standard margins for such a large and mature industry. This shows up noticeably when demand compresses.

Fundamentals: Receivables

Receivables at semiconductor vendors have recently begun to decline from the historical 45-50 days. Receivable days sometimes spike during demand downturns as OEMs conserve cash, stretching payment cycles, but they usually correct after a couple of quarters.

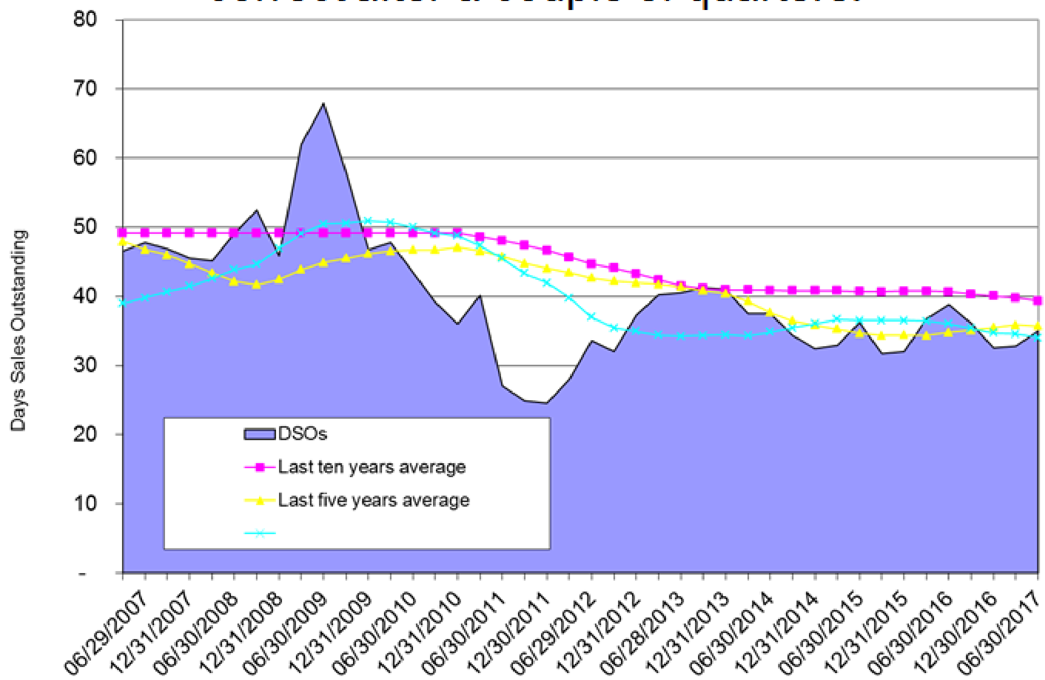

Source: FactSet, Oppenheimer & Co.

Turning to the balance sheet, receivables at semiconductor vendors have historically held at around 45-50 days and are to some degree seasonally affected. During demand downturns, receivables can spike as OEMs look to conserve cash by stretching payment cycles. Receivables have normalized lower since the most recent downturn and have settled in around 35-40 days.

The DSO (days sales outstanding) ratio can also be affected by a quick jump or drop in the sales denominator. Thus, we see DSOs spike when demand falters and fall when it recovers quickly.

Fundamentals: Inventories

OEMs are reluctant to hold inventories except in the tightest supply environment or most robust demand environment. At the same time, the rise in outsourcing allows chip vendors to hold more inventory with less risk.

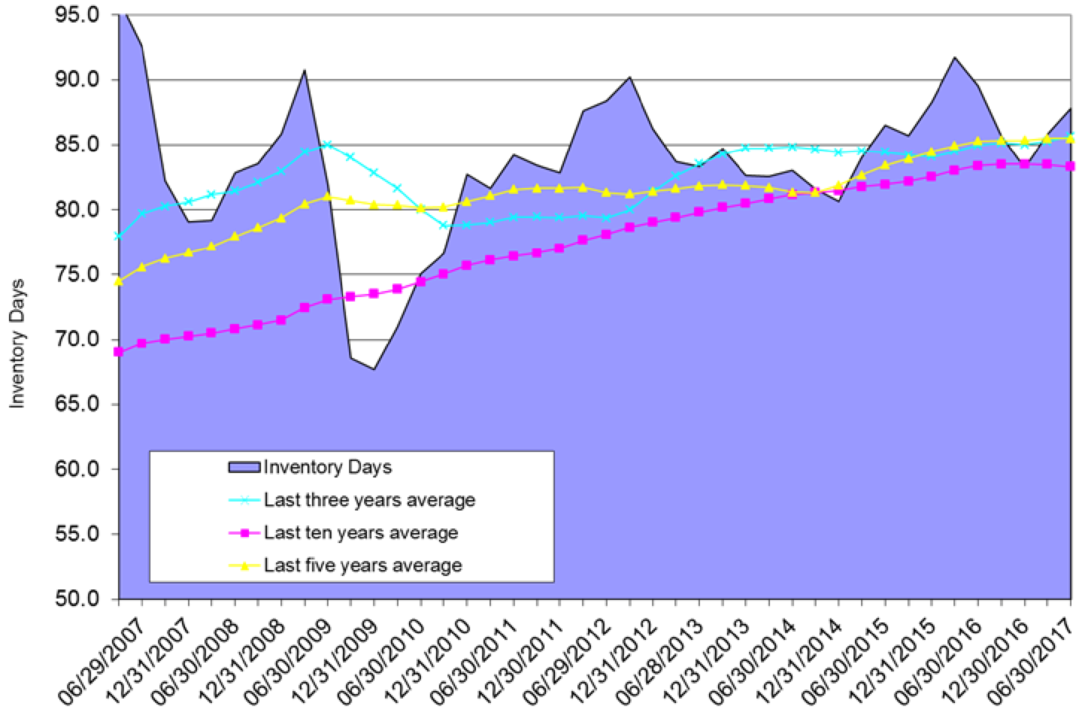

Source: FactSet, Oppenheimer & Co.

Inventory days at semiconductor vendors have varied considerably over the past few cycles. Inventory days have been volatile over the past seven years, spiking above 90 in both 2006 and 2009 while dropping below 70 later that year (2009). Inventories subsequently recovered through mid-2012 before correcting later that year. Inventories declined through mid-2013 as a spike in demand outpaced supply, pushing inventory back within historical ranges. In late 2014, and into 2015, slowing demand steadily pushed inventory levels back up above 85 days, however for the most part they have remained between 80 and 90 days for the last several years.

Semiconductor vendors have, over time, shown increased focus on developing an inventory strategy to meet their customers' desire to minimize risk. IC vendors are holding more inventory on their books on behalf of their OEM customers, in some cases even under a consignment relationship. This may mean relatively high inventory stocks at semiconductor vendors for the foreseeable future. The rise of outsourcing supports this increased inventory with less risk.

Fundamentals: Inventories

Even more important than inventories at semiconductor vendors themselves are inventories downstream, at the OEMs, EMS providers, and distributors. This is the key variable driving the cyclicality of the industry at this point. When these inventories run too low, OEMs tend to double-book in order to guarantee supply to meet demand forecasts. When they run too high, OEMs tend to cancel and push out orders until end demand returns.

PC & Server OEMs

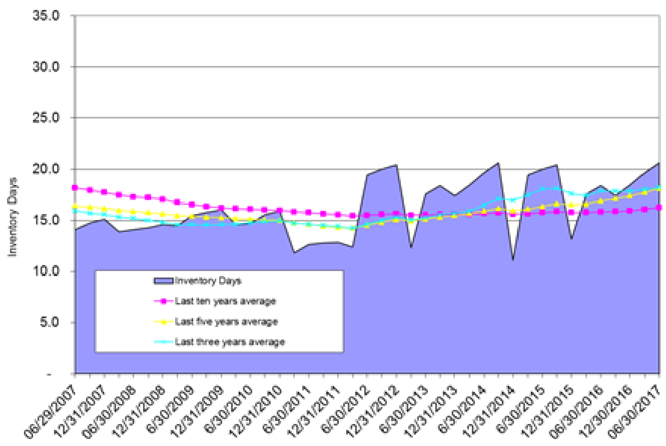

Source: FactSet, Oppenheimer & Co.

Hard Disk Drive OEMs

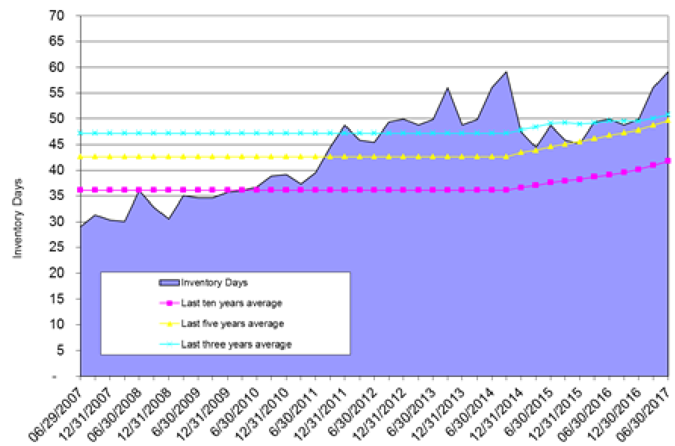

Source: FactSet, Oppenheimer & Co.

Continued…

More important than inventories at semiconductor vendors are inventories downstream, at the OEMs and EMS providers. These are key variables driving the cyclicality of the industry today. When these inventories run too low, OEMs tend to double-book to guarantee that supply meets demand forecasts, lifting sales for semiconductor vendors. When they run too high, OEMs cancel and push out orders until end demand returns, causing semiconductor sales to fall off, often abruptly. We summarize below (and on the next page) the state of inventories by end market:

- *PC and Server OEMs*: Inventory days have normalized to 15-20 days during the past 10 years. This drop has been critical in reducing PC price points and retaining profit margins, especially since the cost of microprocessors and DRAM are two major components of the total bill-of-materials. At present, PC and server hardware OEM inventories are at below-average levels, consistent with an unfavorable and highly uncertain demand environment.
- *Hard Disk Drive (HDD)*: HDD inventory days have consistently risen over the past ten years, as end unit demand (PCs) has stagnated. There has been a considerable amount of consolidation at the OEM level, likely playing a part in this trend as visibility and supply chain management improve. We do not believe this trend of ever-increasing inventories can continue as supply and demand rationalize.

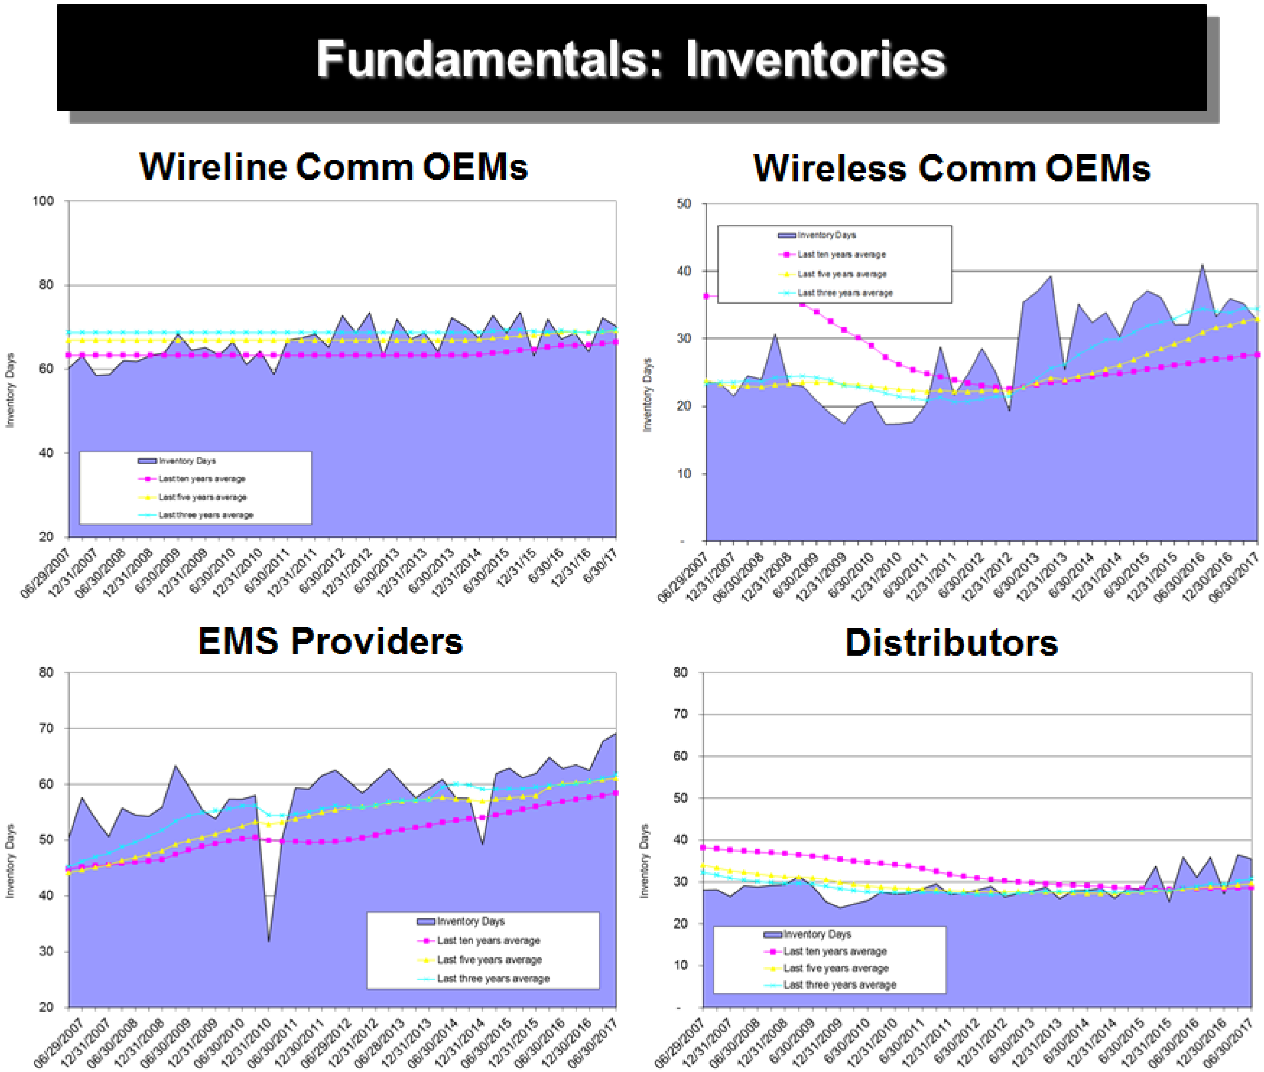

Source: FactSet, Oppenheimer & Co.

- *Wireline OEMs*: Given the long product design and life cycles associated with wireline infrastructure equipment, wireline OEMs have traditionally held higher levels of inventory (50-70 days range) than many tech peers. Inventory levels are high, but managed prudently with little risk of obsolescence. OEMs have adjusted purchasing behavior and have taken advantage of shorter lead times from suppliers. Currently, OEM inventories are moderately high, but well within historical norms.
- *Wireless OEMs*: Inventories at wireless OEMs have historically ranged around 20-40 days and tend to be choppy given the pace of new standard introductions, and there is some seasonality too—inventories usually jump in Q1, stabilize during the year, and decline sharply in Q4. Inventories are now toward the high end of historical ranges given China smartphone demand softness in 1H17, but we expect levels to normalize in the coming quarters.
- *EMS Providers:* EMS inventories have historically ranged around 40-50 days but have been trending higher closer to 60-70 days. Inventory was strong through most of 2004 and 2005, but has been building since that time as OEMs have shifted the burden of carrying inventory to EMS partners. EMS inventory days are now at 68, near a historical high.
- *Distributors:* Distributors build their businesses around efficient inventory management and have been able to limit inventories historically to 25-45 days. Inventory last spiked in 2000/2001 (near 65 days) due to double-ordering. Since then, distributor inventories have been quite healthy, fairly consistent, and currently sit well within historical ranges.

Fundamentals: Debt

Given the long-term nature of manufacturing assets, the industry has maintained some level of debt, with debt/equity ratio fluctuations driven mostly by changes in the equity line. More recently, accelerated buybacks, and large debt issues to fund M&A are becoming more common as organic growth has become more elusive, and interest rates remain at historical lows.

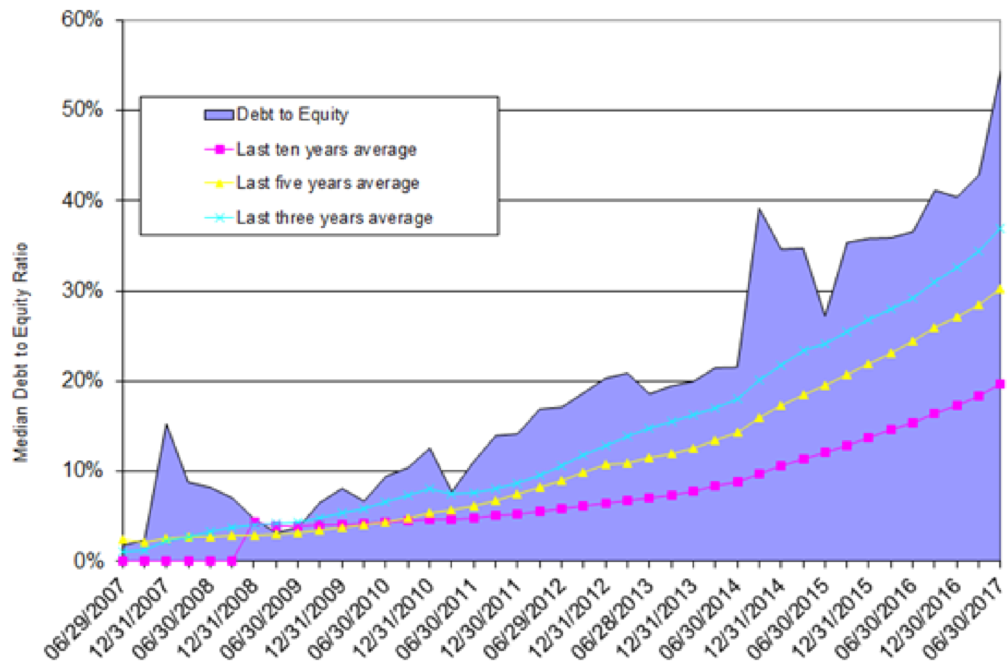

Source: FactSet, Oppenheimer & Co.

Given the long-term nature of semiconductor manufacturing assets, the industry has historically maintained some level of debt, primarily (though not exclusively) concentrated in the hands of IDMs and outsourcers. The fluctuations in debt/equity ratios have tended to be driven more by changes in the equity denominator than any firm trend on the debt side. For example, in 1999 and 2000, the debt/equity ratio penetrated 15% as fabless and IDM semiconductors alike took advantage of inflated multiples in the equity market to raise capital.

As valuations fell, however, the attractiveness of issuing further equity diminished. The ratio was also affected by write-downs of the equity denominator during the 2001 downturn; this included not only inventories and manufacturing assets but also a large amount of intangible assets (primarily acquisition-related), as well as investments in public and private affiliates and even deferred tax assets. The debt/equity ratio thus rose to over 23% in 2003, its highest level in ten years.

In 2007, however, we saw a dramatic change. Buybacks increased across the industry and some companies even recapitalized using debt. This drove the debt/equity line higher. Following the 2008-2009 downturn, semiconductor companies' appetite for debt initially decreased, before accelerating in 2011-14 in an effort to capitalize on record-low interest rates. Debt financing in the low-rate environment has become an attractive way to finance acquisitions and large stock buyback programs. As industry growth slows, many companies have begun to focus on optimizing the capital structure and consolidation. After increasing from $21B in 2014 to $83B in 2015, M&A value increased again to $105B in 2016, led by QCOM's record $46B acquisition of NXPI.

Section III: Market Segments

This section summarizes the major product and end market segments of the industry as well as the competitors within each

Topics include:

Semiconductor Device Types

- Analog vs. Digital: "How a device WORKS"
- SIA Framework: "What a device IS"
- Oppenheimer Framework: "What a device DOES"

Key Competitors

- By Product
- By End Market

Semiconductor Device Types

Analog vs. Digital

"How a device WORKS"

Analog – Analog semiconductor devices deal in precise electric properties, most commonly voltages. Transistors within the device are designed to measure and manipulate these properties. Analog devices are well suited to processing real-world signals, as electronic patterns are used to directly represent the original.

Digital - Digital devices do not deal in the values of actual voltages; rather they simply detect the presence or absence of a voltage. The presence of a voltage is represented digitally as a "1," with the absence represented as a "0." These 1s and 0s can be processed and manipulated digitally with great flexibility.

Mixed Signal – These devices include both analog and digital circuitry. Mixed signal devices are difficult to design and build but can bring the benefits of both analog and digital processing together.

Semiconductor devices are routinely described as "analog" or "digital," which explain how the devices work from a fundamental level. Analog semiconductor devices deal in precise electric properties, most commonly voltages, but also frequency, current, or charge. Transistors within the device are designed to measure and manipulate these properties. Analog devices are well suited to processing real-world signals, as electronic patterns are used to directly represent the original (the word "analog" is related to "analogy"). Since an analog signal has a theoretically infinite resolution, it will always have a higher resolution than any digital system where the resolution is in discrete steps.

Digital devices do not deal in the values of actual voltages (or other properties); rather, they simply detect the presence or absence of a voltage. The presence of a voltage is represented digitally as, a "1," with the absence represented as a "0." The ones and zeroes can be easily manipulated and processed, and since they are universal, they can be used to represent almost anything and perform virtually any electrical function. Often, analog devices will convert a signal to digital (a process known as "quantization") for more intense processing in a digital format—techniques such as compression and error correction cannot be performed in an analog fashion. Certain processes are also inherently digital, such as processing of a computer programming language.

Increasingly, we are seeing semiconductor vendors build mixed-signal devices, which include both analog and digital functionality. These devices are difficult to design and build but can bring the benefits of both analog and digital processing together.

Semiconductor Device Types

The SIA Framework

"What a device IS"

Discretes, Optoelectronics, and Sensors - Includes all non-integrated circuit semiconductor devices. A discrete is a single transistor in a package. Sensors are discrete devices that measure real-world input. Optoelectronics are discrete devices that produce or measure light.

Analog - Devices used to process real-world signals using electronic voltage patterns that represent the original. Includes SLICs (standard linear components) and ASSPs (application-specific analog ICs).

Microcomponents - All digital processors, including microprocessors (MPUs), microcontrollers (MCUs), and digital signal processors (DSPs).

Logic - All non-microcomponent digital logic. Includes ASICs (custom logic), ASSPs (standard specialty logic products), FPGAs (programmable logic), display drivers, and general purpose logic.

Memory - Memory devices are used to store data either for short periods of time or permanently. Includes volatile (DRAM, SRAM) and non-volatile (flash, ROM) memory.

The Semiconductor Industry Association (SIA), the industry body that reports monthly semiconductor industry sales, divides the industry into five broad categories and generally focuses on identifying the type of device sold. When we discuss industry segmentation in terms of the type of semiconductor product, we will generally use the SIA framework.

- **Discretes, Optoelectronics and Sensors:** Includes all non-integrated circuit semiconductor devices. Discretes are devices that contain just a single transistor in a package. These include diodes, small signal and switching transistors, power transistors, rectifiers, thyristors and other discretes.

  Optoelectronics include displays, lamps, couplers, laser devices, image sensors, infrared and other optoelectronics.

  Sensors include those devices that measure temperature, pressure, displacement, velocity, acceleration, stress, strain, or any other physical, chemical, or biological property. The segment also includes actuators, which are responsible for movement.

The SIA Monthly "Blue Book," Which Tracks Monthly Billings Data

- **Analog:** Analog circuits are those used to process real world analog signals, using electronic voltage patterns within the device to represent the original signal (the word analog is related to “analogy”).

  Analog circuits are divided into two broad categories: SLICs and ASSPs. Standard linear integrated circuits (SLICs) are standard analog components used in multiple applications. Though they perform a specific function, they are not tailored to any individual application. SLICs are usually purchased “off the shelf,” often through distribution. Key types are amplifiers, interface circuits, voltage regulators, data converters and comparators.

  ASSPs are more specialized—they are targeted at individual applications (though not individual customers or platforms—that would be classified as a custom ASIC). Analog ASSPs are segmented by application, specifically: automotive, consumer, computer, telecom, industrial and other.

- **Microcomponents:** This includes digital processors, including microprocessors, microcontrollers and DSPs. Digital devices do not deal in the specific voltages of electrons, and instead deal in “1s” and “0s”—“1” being the presence of a voltage and “0” being the absence of a voltage. The ones and zeroes can be used to represent just about anything, as long as they conform to some predefined pattern or code as dictated by software. Note that since digital processors deal only in digital signals, any real world input must be converted to digital before it can be processed by a digital device.

  The microcomponent segment includes the more generic digital ICs—programmable devices that can be used in a variety of applications for computation or signal processing. Microprocessors (CPUs or MPUs) are large, complex processors found in computing and other compute-intensive applications, and typically form the brains of an entire electrical system.

  Microcontrollers (MCUs) are smaller, stand-alone processors that perform dedicated or embedded compute functions within a system. They are usually designed with on-chip RAM, ROM and logic in order to minimize support chip count.

  Digital signal processors (DSPs) are specialized digital processors typically used to process real-time data.

  We note that until 2001, the microcomponents category also included microperipherals, which are dedicated logic devices used to augment MPU or MCU system performance, such as graphics controllers, PC chipsets, Ethernet and modem controllers and hard disk drive controllers. These devices are now classified in the special purpose logic category.

- **Logic:** This segment includes all non-microcomponent digital logic circuits, including gate array and standard cells (ASICs), field programmable logic devices (PLDs or FPLDs), display drivers, general purpose logic and special purpose logic/microperipherals (which are somewhat synonymous with ASSPs).

- **Memory:** Memory devices are used to store data in silicon, either for short periods of time or permanently.

  Volatile memory includes those memory devices that lose stored information when the power is lost. DRAM (Dynamic Random Access Memory) and SRAM (Static Random Access Memory) are the two main types.

  Non-volatile memory retains the stored information when the system is powered off. The two main types are NOR and NAND flash. Older types include mask programmable ROM, EPROM and EEPROM.

Semiconductor Device Types

The Oppenheimer Framework

"What a device DOES"

Analog Processing - Devices used to process real-world signals. Includes SLICs (standard analog components) and ASSPs (application-specific ICs).

Digital Processing - All digital devices used to process signals converted from the analog world. Includes digital signal processors (DSPs), custom and standard digital logic and micro-peripherals, some microcontrollers, and any other digital device not responsible for system control or memory storage.

System Processors - All processors that control the electronic device at the system level, including all microprocessors (MPUs) and most microcontrollers (MCUs).

Memory - All volatile (DRAM, SRAM) and non-volatile (EPROM, flash) memory. Identical to SIA category.

System Enabling Devices – All analog, digital, and discrete devices used as enabling devices on the board. Includes power circuits, glue logic, switches, clocks, discretes, optoelectronics, and sensors.

While the SIA framework is an accurate segmentation of the industry by device type, we often look at the semiconductor industry by end market, with an eye toward system-level design. For this purpose, we prefer to use a simpler, more functional segmentation to help investors understand the various devices and how they fit together in an electrical system.

In this framework, we segment the market into five types of devices, based on the function they perform within the system. Note that we display an icon next to each type of device, which we will consistently use in all block diagrams:

- **Analog Processing:** This includes all analog devices used to process real world signals. Analog ICs deal with complex input from the real world in the form of electrical signals, including sound, light, video, radio waves, temperature, pressure, displacement, velocity, acceleration, stress, strain, or any other physical, chemical, or biological property. They sometimes convert them to digital signals for processing by digital processors or the system CPU. This category would include most of what is found in the SIA product category for analog devices, but would exclude some enabling analog devices such as power management and timing circuits, whose function is secondary to the function of the overall system.

- **Digital Processing:** This includes all digital devices used to perform specialized processing of signals converted from the analog world. This includes DSPs used in wireless and broadband, specialized logic used in datacom and telecom, audio/video processors in consumer electronics, or any other digital processor that does not perform overall system control. Typical functions would be encode/decode, compression/decompression, and acceleration functionality. This category would generally include most of the SIA's MOS logic category (excluding general purpose logic), along with most DSPs and some MCUs.

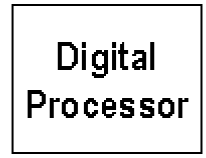

- **System Processors:** This includes processors that perform system control. They are usually highly programmable and run the system software, where applicable. Examples are PC microprocessors, application processors in cellular phones, control plane processors in networking equipment, and microcontrollers that control the majority of system functions. This category would include all microprocessors, most microcontrollers and some DSPs and logic components, depending on how they are used in the system.

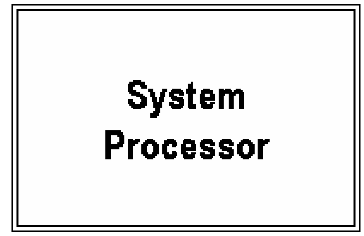

- **Memory:** Identical to the MOS memory category described in the SIA framework, this category includes all volatile and non-volatile memory. Volatile memories like DRAM and SRAM are used by digital processors and system processors, while non-volatile memories are used to store system operating code, boot instructions, and user data.

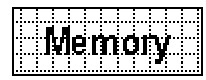

- **System Enabling Devices:** This includes all discrete, analog and digital devices that are used on the board as enabling devices, rather than for processing or storage. Typical functions include switching signals between ICs on the board, stepping up or stepping down voltages, performing system timing and sensing outside input. Examples of enabling devices include discretes, optoelectronics and sensors, as well as power circuits, timing devices, glue logic and some programmable logic devices.

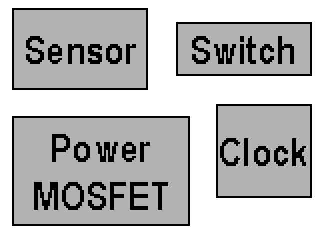

Example: A Basic Electrical System

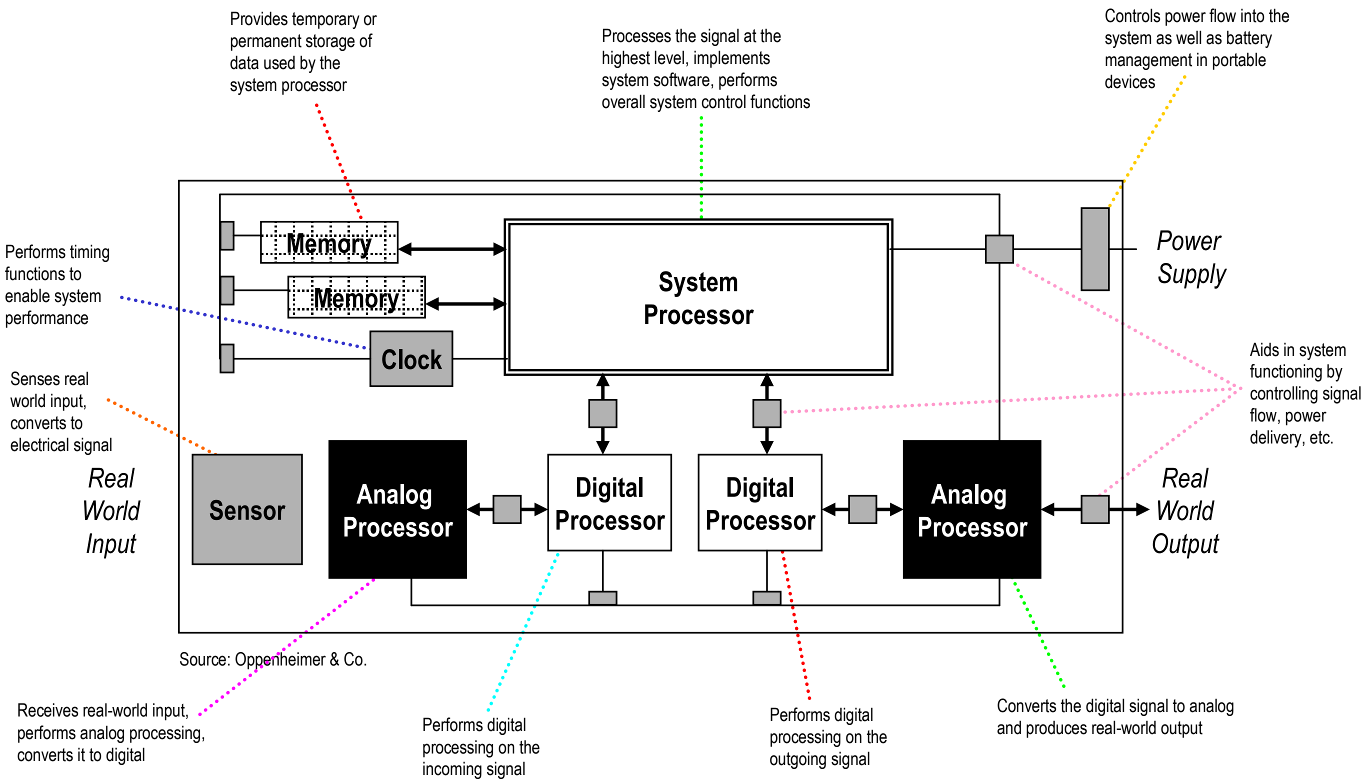

Source: Oppenheimer & Co.

The diagram above depicts an electrical system and identifies the devices on the board using the Oppenheimer framework. In the diagram, some real world input is received, interpreted by a sensor (an enabling device), which converts it to an electrical signal, and is then processed in an analog format by an analog processor. The analog processor converts the signal to digital and passes it to a digital processor, which performs some more intensive digital processing on the signal. The digital processor then passes it to the system processor, which, using software, decides what to do with the signal and can make any changes to it if needed. At some point, the processor sends out the signal to another digital processor, which performs some digital processing and passes the signal on to an analog processor, which converts the signal to analog and produces some real world output.

Devices not directly in the signal chain include memory, which is used for short-term or long-term storage of signals, and other enabling devices, such as power management, discretes, timing devices, switches, display drivers, etc.

A simple example would be a digital camera. When the shutter button is activated, an image sensor (enabling device) receives the visual input, which is processed by an analog front end (analog processor) and converted to a digital format. A DSP or an ASIC (digital processor) then encodes and compresses the signal and tells a microcontroller or microprocessor (system processor) that a picture has arrived. The picture file is then stored in flash memory (memory), and the microcontroller tells an LCD controller (digital processor) to produce an image on the LCD screen. The LCD controller then directs signals from row and column drivers to produce the image on the screen. The table on the following page displays additional examples.

End Market Application	Analog Processors	Digital Processors	System Processor	Memory	System Enabling Devices
Computing					
Desktop PC	Graphics interface, sound interface, Ethernet PHY, modem line driver	Chipset, graphics processor, sound processor	Microprocessor	DRAM, ROM, flash	Power ICs, clocks, switches, discretes
Notebook PC	Graphic/sound interface, Ethernet PHY, modem, WLAN/Bluetooth radio, power management	Chipset, graphics processor, sound processor, WLAN/Bluetooth baseband/MAC	Microprocessor	DRAM, ROM, flash	Power ICs, clocks, switches, discretes, LCD drivers
Hard disk drive	Read channel, preamp, motor driver	Servo ASIC, hard disk controller	Microcontroller (embedded in hard disk controller)	DRAM, SRAM, flash	Magnetic sensors, power ICs, clocks, switches, discretes
Communications					
Wireless handset	Power amp, radio (RF/IF), audio codec, Bluetooth radio, analog baseband	Digital baseband (DSP), Bluetooth baseband/MAC, Camera image processor, Multimedia processor	Applications processor (can be embedded in DSP)	Flash, SRAM, low-power DRAM	Sensors, power ICs, clocks, switches, discretes, LCD drivers
Ethernet switch	Ethernet PHY	Ethernet MAC, switch chipset, packet processor, programmable logic	MIPS/PowerPC processor	SRAM, DRAM, ROM	Power ICs, clocks, switches, discretes
Wireless LAN access point	Radio (RF/IF), power amp	Baseband/MAC	MIPS/ARM/PowerPC processor	DRAM, SRAM	Power ICs, clocks, switches, discretes
SONET/SDH line card	PMD, CDR, SONET transceiver/PHY, backplane SerDes	Framer, FEC, network processor, traffic manager	MIPS/PowerPC processor	DRAM, SRAM	Power ICs, clocks, switches, discretes
Cable/DSL modem	Tuner/line driver, analog front end	DOCSIS MAC/processor (cable modem), data pump DSP (DSL)	Communications processor	DRAM, SRAM, flash	Sensors, power ICs, clocks, switches, discretes
Fibre Channel Host Bus Adapter (HBA)	Fibre Channel SerDes/PHY	Fibre Channel controller ASIC	Microprocessor (can be integrated into FC controller)	SRAM, flash	Power ICs, clocks, switches, discretes
Consumer					
Digital set-top box/DVR	Tuner, analog front end (mux/demux)	MPEG decoder/encoder, graphics processor	Microcontroller	DRAM, flash, ROM	Power ICs, clocks, switches, discretes
Digital still camera	A/D converter, Bluetooth/ WLAN radio	Digital image processor, ASICs, Bluetooth/WLAN baseband/MAC	Microcontroller, Microprocessor	Flash, DRAM	CCD/CMOS image sensor, power ICs, clocks, switches, discretes, LCD drivers
PMP	Audio driver, D/A converter	DSP, audio/video processor/ASSP	Microcontroller	Flash	Power ICs, clocks, switches, discretes, LCD drivers
Video game console	Audio DAC, video DAC, Bluetooth/WLAN radio, Ethernet PHY	Graphics processor, system I/O, Bluetooth/WLAN baseband/MAC, Ethernet MAC	Microprocessor	DRAM, SRAM, flash	Sensors, power ICs, clocks, switches, discretes, LCD drivers (handheld)

Source: Oppenheimer & Co.

Key Semiconductor Competitors

Semiconductor Market Top 10 Competitors

Intel (US)—The largest semiconductor manufacturer in the world for 25 consecutive years, Intel held 16% of the market in 2014. Intel dominates the PC (client and data center) microprocessor segment and is also a top supplier of PC chipsets, embedded microprocessors, and PC networking components. The company is diversifying into artificial intelligence, handset modems, and NAND flash memory.

Samsung (Korea)—The diversified electronics giant is No. 1 in DRAM and NAND flash. It also supplies microcontroller and logic devices, as well as its own Exynos application processors used in smartphones. Samsung also operates as a foundry for Apple and recently won share from Qualcomm supplying LTE modems to Apple.

Qualcomm (US)—Qualcomm is the leading supplier of handset ICs. It has a No. 1 position in the baseband market and a dominant position in the 4G LTE market. The company's revenue has declined during the last two years as it has been in the process of pivoting its QCT business to markets outside of smartphones as the end market's growth has decelerated. To aid its diversification efforts, Qualcomm announced plans to acquire NXP in October 2016.

SK Hynix (Korea)—Hynix is the No. 2 supplier of DRAM, which accounts for nearly 75% of the company's sales. Hynix is building its position in NAND. 2016 revenue declined 10% as both the NAND and DRAM markets were in oversupply, a dynamic which reversed by mid-2016.

Broadcom (US)— Broadcom has moved up in the rankings due to its acquisition of Avago, which closed in February 2016. Broadcom has focused on enhancing its enterprise and datacenter strength providing network infrastructure and switching products, many of which were through acquisitions. Broadcom also realizes ~30% of its revenue from wireless application through its in connectivity and combo chips (WLAN, BT, GPS, FM, NFC), and legacy Avago's FBAR RF filters.

Micron Technology (Japan)—Micron is a top supplier of DRAM and NAND and won the bidding process in mid-2012 to acquire bankrupt Elpida Memory, also a leader in DRAM and NAND. Micron is working on a new type of memory, 3D XPoint, in a joint effort with Intel, targeting the datacenter market.

Texas Instruments (US)—The top supplier of analog ICs and DSPs. TI is focused on the analog and embedded markets and exited its wireless/connectivity business. The company sells products to a broad array of end markets, most notably the auto and industrial markets which account for roughly 50% of revenue and have been an area of investment for the company due to the end markets sticky and stable product-cycles, and content growth opportunities.

Toshiba (Japan)—Toshiba is the No. 2 supplier of NAND, and also a major supplier of ASICs, optoelectronics, and standard logic. Due to financial issues in Toshiba's other businesses, however, the company announced plans to sell its memory business to an international consortium.

NXP Semiconductor – NXP moved up to the No. 9 position in 2016 due to its acquisition of Freescale. NXP supplies a range of products including MCUs, power-management ICs, interface, and connectivity ICs. NXP is the leading semiconductor supplier to the automotive market. The company announced plans to be acquired by Qualcom in October 2016.

MediaTek – MediaTek grew its revenue 30% organically in 2016 by refilling channel inventories in China and abroad, supplying products to mid-tier and high-end consumer devices, most notably handsets. Its SoCs serve as a less expensive offering to Qualcomm's products. MediaTek also completed an acquisition of Richtek, a supplier of power management ICs, boosting revenue ~7%.

Key Semiconductor Competitors

Semiconductor Market Top Competitors By Product

Analog

Analog SLICs		Analog ASSPs	
Texas Instruments	Cirrus Logic	STMicro	Qualcomm
Analog Devices	STMicro	Infineon	Qorvo
Maxim	Richtek	Texas Instruments	Broadcom
ON Semi	Sanken	Renesas	Skyworks
		NXP	

Microcomponents

Microprocessors		Microcontrollers	DSPs
PC & Server	Embedded		
Intel	Intel	Renesas	Texas Instruments
AMD	NXP	Freescale	Analog Devices
NXP	Cavium	STMicro	NXP
	Texas Instruments	Microchip	Cirrus Logic
	Broadcom	NXP	Toshiba
	AMD	Texas Instruments	
		Infineon	
		Amtel	
		Cypress	
		Samsung	

Digital Logic

ASICs	PLD	GP Logic	Display Drivers
AMD	Xilinx	Xilinix	Samsung
Samsung	Intel	Intel	Novatek
Broadcom	Microsemi	Samsung	Himax
IBM		Novatek	Renesas
STMicro		Himax Technology	Ilitek
Dialog Semiconductor		Texas Instruments	Orise Tech
Renesas		Renesas	Silicon Works
Texas Instruments		NXP	Magnachip
Shenzhen HiSilicon		Ilitek	Raydium
			Synaptics

Memory

DRAM	SRAM	Flash	
		NOR	NAND
Samsung	Cypress	Cypress	Samsung
SK Hynix	Renesas	Macronix	Toshiba
Micron	GSI Technology	Winbond	Western Digital
Nanya	Microchip	Samsung	Micron
		Microchip	SK Hynix
			Intel

Mask ROM	EPROM	EEPROM	
Macronix	Rohm	STMicro	Rohm
	Microchip	Microchip	Microchip
	Tongfang Guoxin	NXP	ON Semi
		Infineon	
		Seiko	

Discretes and Optoelectronics

Discretes		Sensors	
Infineon	Vishay	Robert Bosch	InvenSense
NXP	Toshiba	STMicro	Denso
Mitsubishi	Rohm	Infineon	Knowles
ON Semi	Renesas	Asahi Kasei	Analog Devices
STMicro		Sanken	

Optoelectronics			
Image Sensors		All Other Optoelectronics	
Sony	Toshiba	Broadcom	Microchip
Samsung	Sharp	Vishay	ON Semi
ON Semi	SK Hynix	Osram	Sharp
	STMicro	Renesas	
		Mellanox	

Source: Dataquest, IDC, Forward Concepts, Databeans, Oppenheimer & Co.

Key Semiconductor Competitors

Semiconductor Market Top Competitors By End Market

Computing

PCs

CPU	PC DRAM	Graphics	Clocks/Buffers	I/O Controller	Analog Modem	PC Ethernet	PC Wireless LAN	Power Mgmt
Intel	Samsung	NVIDIA	IDT	Standard Micro	Synaptics	Broadcom	Intel	Texas Instruments
AMD	Hynix	AMD	Intel	Winbond/National	Silicon Labs	Intel	Qualcomm Atheros	Maxim
Via	Micron		Realtek	ITE	Broadcom	Realtek	Broadcom	Analog Devices
	Nanya		Cypress			Marvell		STMicro
	Powerchip		Maxim	**Audio**		Qualcomm Atheros		ON Semi
	Winbond		Semtech	Creative				Sanken
Chipsets				Realtek			**PC Bluetooth**	Richtek
Intel				IDT			Qualcomm - CSR	Power Integrations
AMD				Conexant			Broadcom	Renesas
							Microchip	

Computing

PC Displays		Hard Disk Drives	Optical Drives		Printers/MFPs
LCD Controllers	**LCD Drivers**	Marvell	**Controllers**	**Optoelectronics**	STMicro
MediaTek/M-Star	Samsung	Broadcom	MediaTek	Sharp	Texas Instruments
Realtek	Novatek	Texas Instruments	Renesas	Broadcom	Marvell
Novatek	Himax	STMicro	Panasonic	Stanley Electric	Synaptics
	Renesas	Renesas	NEC	Sony	Microsemi
	Ilitek		Ricoh	Toshiba	Toshiba
	Orise		Rohm		
	Silicon Works				
	Magnachip				
	Raydium				
	Synaptics				

Networking

Ethernet			Wireless LAN	Bluetooth	Enterprise Storage		
PC Controllers	**PHY & Switch ASSPs**	**ASICs**	Qualcomm	CSR	**HBA/SAN ASICs**	**RAID Processors/ Controllers**	**Disk Array ICs**
Broadcom	Broadcom	Broadcom	Broadcom	Broadcom	Broadcom	Intel	Microsemi
Intel	Marvell	IBM	Marvell	Microchip	IBM	Broadcom	Broadcom
Realtek	Intel	TI	Realtek	Qualcomm		NXP	IBM
Marvell	Realtek	Fujitsu	Mediatek	Nordic Semi		Marvell	Maxim
Microsemi	Qualcomm			Marvell		Microsemi	

Telecom/Datacom

CO Line Cards	Analog Modem	DSL	Cable Modem	PON	VoIP
Infineon	Synaptics	Broadcom	Broadcom	Broadcom	Texas Instruments
Zarlink	Silicon Labs	Intel	Texas Instruments	Microsemi	Silicon Labs
STMicro	Broadcom	Qualcom	Qualcomm - CSR	Inphi	Broadcom
IDT	Analog Devices	MediaTek		Marvell	Microsemi
NXP				Qualcomm	MACOM
				Semtech	Infineon
				MACOM	Broadcom
					NXP

Telecom/Datacom

Communications Infrastructure

T/E	SONET/SDH	ATM	Comm Processor	Network Processing: ASICs	Network Processing: NPUs	Network Processing: CAM/NSEs	Other	Switching: ASICs	Switching: ASSPs
Infineon	MACOM	Microsemi	NXP	IBM	Intel	Broadcom	Cavium	IBM	Broadcom
Microsemi	Inphi	MACOM	Intel	Hisilicon (Huawei)	Marvell	Renesas	Maxlinear	Fujitsu	MACOM
Maxim	Broadcom	Inphi	Cavium	Broadcom	Microsemi		Gemalto	Microsemi	Microsemi
MACOM	Vitesse		Broadcom	Texas Instruments	Mellanox		Broadcom		
Broadcom	Microsemi			Global UniChip	Broadcom		Microsemi	**Interconnect**	
MaxLinear				STMicro				Broadcom	
NXP				Broadcom				Intel	
								IDTI	
								Diodes	

Key Semiconductor Competitors

Semiconductor Market Top Competitors By End Market

Wireless

Handset ICs					Infrastructure			
Digital Baseband	**Analog Baseband**	**App. Processor**	**RF & PA**	**Memory**	**Baseband ASIC**	**Baseband DSPs**	**Transceiver/ Converter**	**Power Amp**
Qualcomm	Qualcomm	Apple	Qualcomm	Samsung	Qualcomm	NXP Semi	Analog Devices	Qorvo
Mediatek	Maxim	Samsung	Broadcom	Hynix	Texas Instruments	Analog Devices	Texas Instruments	Skyworks
Intel	Cirrus	Intel	Skyworks	Toshiba	Broadcom	Texas Instruments	KEMET	Anadigics
Spreadtrum	Dialog		MediaTek	Micron			IBM	NXP
RDA	Samsung		Intel	Western Digital			Maxim	
Samsung	Broadcom		Qorvo	Nanya				
				Macronix				

Digital Consumer

Set-Top Box	Digital TV	DVD	Digital Camera	Videogame	Memory Card	
			Image Sensor		**NAND**	**Controller**
STMicro	Samsung	Mediatek	Sony	Renesas	Samsung	Kingston
Broadcom	NXP	Broadcom	Sharp	IBM	Toshiba	Toshiba
MaxLinear	Toshiba	Sigma Designs	ON	Sony	Western Digita	Samsung
Intel	Panasonic	Renesas	Toshiba	Sharp	Micron	Transcend
Sony	Himax	Sony	Samsung	Samsung	Hynix	Micron
Samsung	Mediatek	Samsung	Panasonic	NVIDIA	Intel	Marvell
	Sony	Panasonic	E2V	AMD		
	STMicro			Marvell		
	Broadcom			Qualcomm Atheros		
	MaxLinear			Broadcom		

Automotive

Automotive			
Auto MCUs	**ASICs/ASSPs**	**Auto SLICs**	**Auto Discretes**
Renesas	STMicro	Texas Instruments	Infineon
NXP	Infineon	Analog Devices	Renesas
Texas Instruments	NXP Semi	Maxim	NXP
STMicro	Texas Instruments	ON Semi	Rohm
Infineon	Robert Bosh	Sanken	ON Semi
Microchip	Renesas	Toshiba	Mitsubishi
Cypress	Panasonic	Renesas	Toyota
Denso	Toshiba	Infineon	Denso
Toshiba	Denso		Toshiba

Source: IDC, Dataquest, Oppenheimer & Co.

Section IV: Product Summary

This section discusses the major product segments of the semiconductor market

For each product we discuss:

Product Definitions

Market Forecast

Competitive Environment

The Semiconductor Market

Semiconductor Industry Revenue Forecast

•The semiconductor industry totaled $339 billion in sales in 2016, up 1% from 2015
•Revenue grew 1% in 2016, with a weak memory pricing environment muting growth. Pricing has strengthened significantly in 2017
•We are expecting moderate growth going forward, overall, we expect a 5-year CAGR of 5%, with sales increasing from $339 billion in 2016 to $426 billion in 2021.

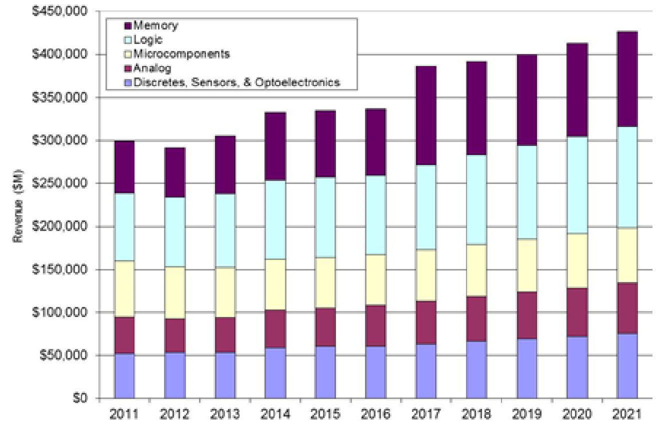

Source: WSTS, Oppenheimer & Co.

Product Forecast Highlights

•Microcomponents should see modest growth, as MCU unit growth of ~8% is partially offset by an expected ~6% drop in APSs.

•Analog should grow at a 4% CAGR with near equal growth in standard linear ICs and application-specific ICs

•Logic should grow slightly below 5%. Display drivers are expected to grow below the market, while special purpose logic and FPGAs grow above.

•Memory should significantly outperform after a banner 2017 which could grow 50%+ on a strong pricing environment. Growth should moderate thereafter off the high base and likely remains volatile with swings in supply/demand balance.

•Despite perpetual price pressure, discretes, sensors, and optoelectronics should slightly below the overall industry CAGR given their multi-market application at just over 4%

The semiconductor industry totaled $339 billion in sales in 2016, up 1% from 2015. Revenue grew 1% in 2015, with a weak memory pricing environment muting growth. Pricing has strengthened significantly in 2017. We are expecting moderate growth going forward; overall, we expect a 5-year CAGR of 5%, with sales increasing from $339 billion in 2016 to $423 billion in 2021. Microcomponents should see modest growth, as MCU unit growth of ~10% is partially offset by an expected ~7% drop in ASPs.

- Analog should grow at a 4% CAGR with near equal growth in standard linear ICs and application-specific ICs
- Logic should grow slightly below 5%. Display drivers are expected to grow below the market, while special purpose logic and FPGAs grow above.
- Memory should significantly outperform after a banner 2017 which could grow 50%+ on a strong pricing environment. Growth should moderate thereafter off the high base and likely remains volatile with swings in supply/demand balance.
- Despite perpetual price pressure, discretes, sensors, and optoelectronics should slightly below the overall industry CAGR given their multi-market application at just over 4%

Product Summary

The semiconductor market can be divided into five key product areas

Five Key Markets:

Analog

Microcomponents

Digital Logic

Memory

Discretes and Optoelectronics

We first examine the semiconductor industry by product group. The industry can be divided into five segments:

- **Analog:** Analog circuits are used to process real world signals, using electronic voltage patterns within the device to represent the original signal. Standard linear integrated circuits (SLICs) are standard analog components used in multiple applications, usually purchased “off the shelf,” often through distribution. Analog application specific standard products (ASSPs) are more specialized analog devices targeted for a specific application.
- **Microcomponents:** These programmable devices perform intensive compute processing and system control. Microcomponents include microprocessors, microcontrollers and DSPs.
- **Digital Logic:** Logic devices perform specialized digital processing, and include gate arrays and standard cells (ASICs), programmable logic (PLDs, FPGAs), display drivers, general purpose logic and special purpose logic.
- **Memory:** Memory devices are used to store data either for short periods of time or permanently. Volatile memory loses stored information when the power is lost; key types are DRAM and SRAM. Non-volatile memory retains the information when the system is powered off; key types are flash, mask ROM, EPROM and EEPROM.
- **Discretes, Optoelectronics and Sensors:** This category includes all non-integrated circuit semiconductor devices. Discretes contain just a single transistor in a package. Optoelectronics are specialized discretes designed to emit and detect light. Sensors measure physical, chemical or biological properties.

Analog

Analog:

Devices designed to process, convert, and create analog real world signals

Two Key Markets:

Analog SLICs

Analog ASSPs

Analog ICs are devices used to process, convert, and create real world signals. Whereas digital circuits operate by determining the absence or presence of an electrical charge (the "1s" and "0s" of the digital world), analog circuits are concerned with the actual voltages (or other electrical properties) in circuits. When electronic systems interact with the real world, analog circuits represent this real world input in the form of voltage/current strength (the electrical charges in the device form an "analogy" of the real world input—hence the term "analog"). They can also convert the signal to a digital form for processing. Conversely, when the device creates real world output, analog circuits convert the digital instructions to analog voltages/currents to create an electrical output.

Analog circuits are divided into two broad categories: **SLICs** and **ASSPs**.

- **SLICs** (standard linear integrated circuits) are standard analog components that are used in multiple applications. Although designed for a specific function within a system, a SLIC is not optimized for particular types of systems. Key types are amplifiers, interface circuits, voltage regulators, data converters and comparators.
- **ASSPs** (application specific standard products) are more specialized analog devices targeted for a specific application. These devices are often designed as multiple SLICs that are integrated and then modified to serve a specific function in a particular end market; but they can also be designed from the ground up. Examples include radio frequency (RF) ICs in cellular phones, interface ICs for networking, read channels in HDDs, tuners and demodulators in set-top boxes, battery management in notebooks and power train control devices in automobiles.

Analog SLICs

Analog SLICs (Standard Linear Integrated Circuits) - Standard analog components used in multiple applications. While designed to perform a specific function, SLICs are not optimized for any particular type of system.

Key Types: Amplifiers, interface drivers, voltage regulators, data converters (analog-to-digital, digital-to-analog, codec), comparators.

Market Forecast

- Analog SLICs are expected to grow at a 5-year CAGR of 5%, rising from $19.5 billion in 2016 to $24.4 billion in 2019.
- Voltage regulators should lead growth, with interface, data converters and amplifiers all slightly trailing.

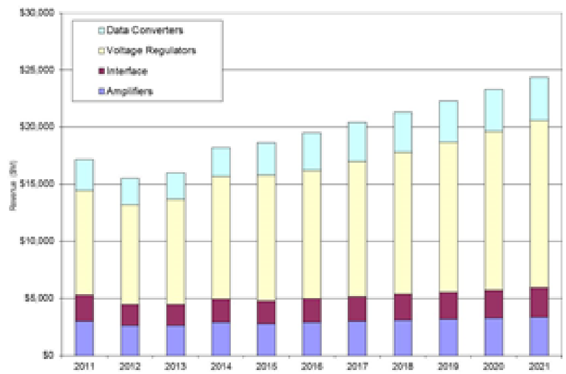

Source: WSTS, Oppenheimer & Co.

Key Competitors

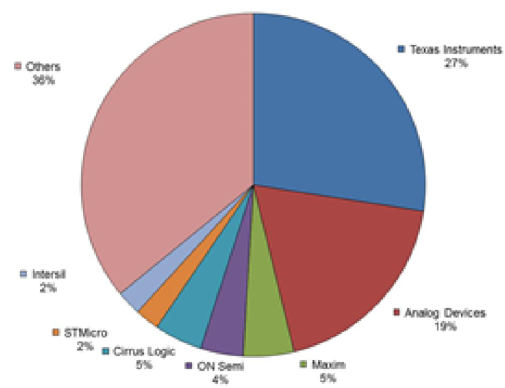

Source: Dataquest, Oppenheimer & Co.

Analog SLICs (standard linear integrated circuits) are general purpose analog components that can be used across multiple types of electronic systems. Although designed to perform a specific function within a system, SLICs are not optimized for particular applications. For example, a line driver (a type of interface SLIC) is designed to produce or receive a signal on a copper wire, but can be used in any application that requires this function (e.g., a telephone set, an analog modem). Similarly, a video ADC (analog-to-digital converter) is designed to convert analog video signals to digital, but can be used in any video application (e.g., digital camera, digital TV, set-top box).

SLICs can be sub-divided into five types of devices:

- **Amplifiers**: ICs designed to amplify a signal, also can be configured to perform certain related operations.
- **Interface**: ICs that produce output or receive input in the transmission of signals to and from systems.
- **Voltage Regulators**: Devices that control power flow by providing a stable flow of current at a particular voltage.
- **Data Converters**: Includes “mixed signal” (analog/digital) devices designed to convert a signal from analog to digital (an ADC) or from digital to analog (a DAC). A codec combines the two functions in a single device.
- **Comparators**: ICs used to compare two quantities, usually used in power management. This category of device is often grouped with amplifiers, as it is a small market and is similar in nature to amplifiers.

Market Forecast: The market for analog SLICs is expected to grow at a 5% CAGR over the next five years. While most electronics OEMs and semiconductor companies are focused on driving the performance of the digital portion of the system, most applications still require significant analog content to interface with the real world. In total, we expect the market to grow at a 5-year CAGR of 5%, rising from $19.5 billion in 2016 to $24.4 billion in 2019. Voltage regulators should lead growth at a 5% CAGR, while amplifiers and interface drivers are expected to grow at a 4% CAGR and data converters should lag with a CAGR of 3%.

Competitive Environment: The SLIC market is one of the most stable and profitable segments of the semiconductor industry. Barriers to entry in the standard linear market remain extremely high, due in part to the nature of analog SLIC device design. Analog engineering is more of a refined art than a school-taught competency, with most analog engineers honing their craft over years of real world design. While digital devices are more dependent on scaling processing architectures and migrating to smaller geometries, SLICs are more dependent on a deep knowledge of the intricacies of the particular medium for which the device is being designed (e.g., sound waves, light waves, heat). The pool of analog design engineering talent is therefore more limited than for digital devices, and a multitude of SLIC vendors have been able to develop defensible franchises. Further, as SLICs are extremely diverse in application, the supplier chain is extremely fragmented.

Only two vendors held market share of 10% or better in 2014 in the analog SLIC segment. Texas Instruments held 27%, with a leadership position in amplifiers, voltage regulators and interface ICs and a strong No. 2 position in converters. The company acquired perennial share loser National Semiconductor in 2011 to increase its exposure to the Industrial vertical. Analog Devices held 19% after acquiring Linear Technologies in 2017, with a dominant 38% share of data converters and a strong amplifier position. Maxim, the high-performance standard linear pure-play, held a 5% share, and was No. 2 in interface ICs, and No. 4 in voltage regulators. Rounding out the top 5 were ON Semi and Cirrus Logic with 4% and 5% share each We provide additional competitive detail in the charts below.

Amplifiers and Comparators

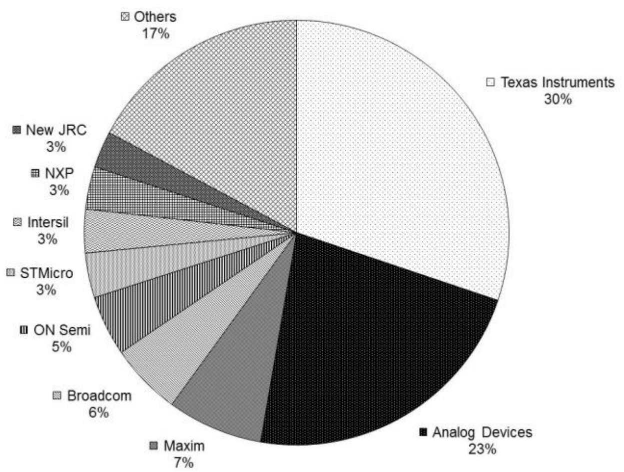

Voltage Regulators

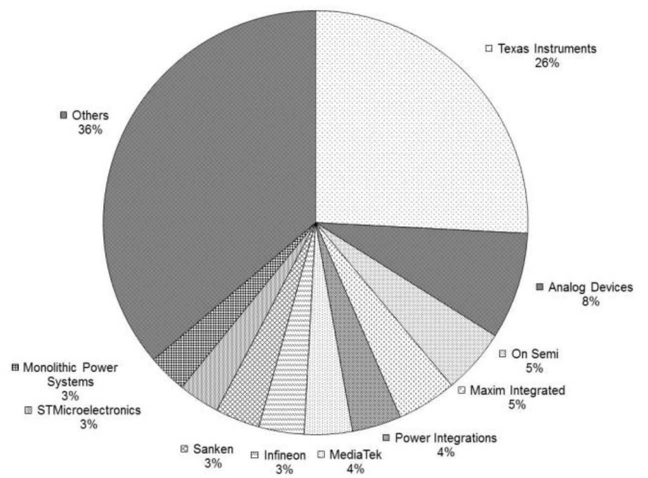

Converters

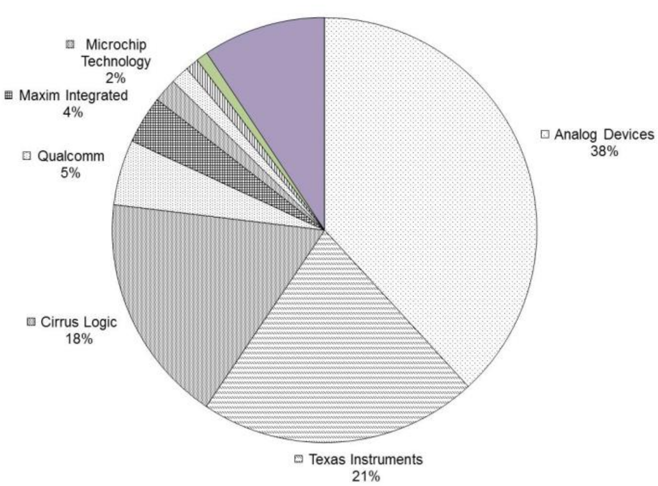

Interface Circuits

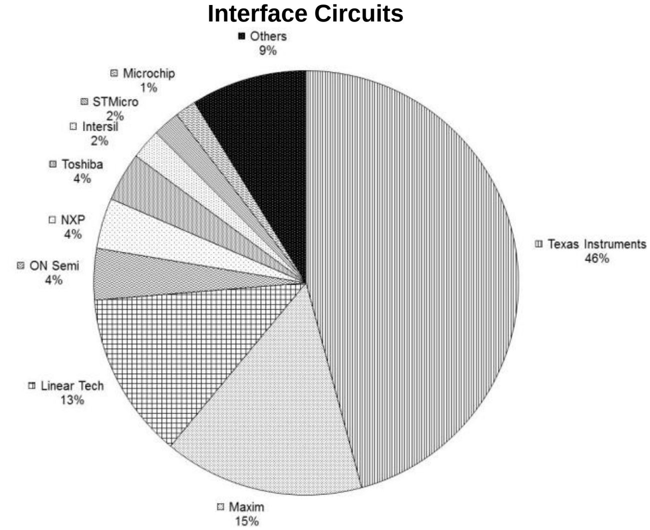

Source: Gartner, Oppenheimer & Co.

Analog ASSPs

Analog ASSPs (Application Specific Standard Products) - Specialized analog ICs targeted for a specific application. Often integrate multiple SLICs or even some digital circuitry. Often called "mixed signal" analog.
Key Types: Computer, communications, consumer, automotive, industrial.

Market Forecast

•Analog ASSPs are expected to grow at a 4% CAGR, rising from $28.3 billion in 2016 to $34.4 billion in 2021.
•The trend toward integration in the digital logic segment continues to weigh on application-specific analog, particularly in the Consumer and Computing segments.
•Growth to be led by Automotive, Industrial and Communications ASSPs.

Market Forecast

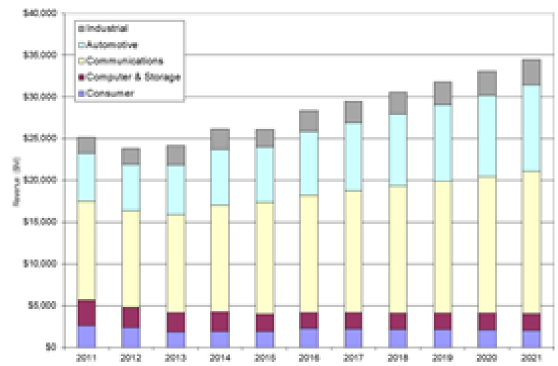

Source: WSTS, Oppenheimer & Co.

Analog ASSPs (application specific standard products) are analog components designed to service a particular application. These are sometimes designed as multiple SLICs that are integrated to serve a specific function in a particular end market, but can also be designed from the ground up. They often include some digital circuitry, and hence are sometimes referred to as "mixed signal analog."

Key examples of analog ASSPs include radio frequency (RF) ICs in cellular phones, physical layer interface devices for networking applications, read channels in hard disk drives, demodulators in consumer electronics, battery management ICs in notebook computers, and engine and power train control devices in automobiles.

Unlike analog SLICs, analog ASSPs are not usually sold through the distribution channel, but instead go through a design process at OEMs that is similar to the one used for digital ASSPs. In fact, many of the highest volume electronics applications are served by a chipset that includes one or more analog ASSPs paired with one or more digital ASSPs offered by the same vendor, or by multiple vendors on a single reference design. The analog and digital ASSPs are thus designed in, and subsequently purchased by the OEM, at the same time.

As mentioned above, some analog ASSPs actually incorporate mixed signal and digital functionality, often blurring the line between analog ASSP and special purpose logic. Further, the same vendors who design digital ASSPs are now increasingly the vendors who design analog ASSPs. While the SIA has for now kept the line between analog and digital ASSPs, some third-party research firms like Gartner have done away with this distinction and group all application specific analog and digital devices as simply "ASSPs."

Analog ASSPs are divided into five categories by application:

- **Automotive**: ASSPs designed for use in automobile engine, safety, comfort, entertainment, power controls, and advanced drivers assistance/autonomous driving
- **Consumer**: ASSPs used in audio/video equipment and home appliances, usually for production, reception, or processing of audio and video input and output.
- **Computer**: ASSPs used in computing applications, most notably storage devices like hard disk and optical drives.
- **Communications**: ASSPs used in wireless, telecom, and datacom client and service provider devices and equipment. Often used in the production and reception of transmission signals (i.e., optical, electrical and radio).

- **Industrial and Other**: ASSPs used in industrial equipment and other applications.

Market Forecast: The market for analog ASSPs is expected to grow at a 4% CAGR, rising from $28.3 billion in 2016 to $34.4 billion in 2019. Analog ASSPs modestly underperform the overall semiconductor industry growth given integration pressure from logic-based SoCs (system-on-a-chip), which limit the TAM for discrete analog components. A prime example of this is in the communications market (49% of total ASSP market), which is currently experiencing the most significant drive toward integration. Wireless handsets, a large consumer of RF, analog baseband and power management ICs, are starting to utilize single-chip solutions that integrate analog functions into digital SoCs. In total, we estimate communications ASSPs should be growing at 5% through 2021 (note that outside of wireless, certain networking technologies should drive some communication ASSP growth, such as 10G/25G Ethernet, Fibre Channel, etc.). PC and consumer are markets that are expected to underperform, with Consumer declining at a 2% CAGR and PC's roughly flat over the same period. In a sluggish consumer spending environment, the lack of "must-have" devices is driving inconsistent unit growth that is also subject to more pricing pressures as well as the trend to pull analog functions into the digital SoC. Computing has already undergone an integration cycle, specifically in HDDs, and so should see a growth rate closer to that of units going forward.

We expect more consistent growth from the industrial and automotive segments, increasing at CAGR of 6% and 7%, respectively. Both industrial and automotive benefit from the increasing complexity of electrical systems and a mix shift towards systems with greater semiconductor content; there is also less integration pressure as FPGAs are the main digital IC in an industrial stem, as opposed to a digital SoC.

	Communications	Computing	Consumer	Auto/Industrial
Key device types	Analog basebands Handset power management RF transceivers Power amplifiers	Printer front ends and cartridge heads HDD read channels, preamps, and motor controllers PC interface ICs	Demodulators Tuners Audio/video interface Digital camera front end DVD front end	Engine control Automotive audio Climate control Engine Control (ECM) Infotainment ADAS
Key Competitors	STMicroelectronics Qualcomm/NXP Qorvo Texas Instruments Intel Renesas NXP Broadcom	Texas Instruments STMicroelectronics MediaTek Renesas Rohm Marvell Broadcom NXP	Sony Samsung NXP STMicroelectronics Renesas Broadcom Texas Instruments	STMicroelectronics NXP Infineon Renesas Robert Bosch Freescale Texas Instruments Toshiba

Source: Gartner, Oppenheimer & Co.

Microcomponents

Microcomponents:

Devices designed to perform intensive compute processing and system control

Three Key Markets:

Microprocessors

Microcontrollers

Digital Signal Processors

Microcomponents include all digital processors, including microprocessors, microcontrollers and DSPs (digital signal processors). These devices are built around the famous ones and zeroes of the digital world and are designed to be programmable; i.e., they run software that is very customizable and can change the function of the device.

Microprocessors (CPUs or MPUs) are the largest, most complex processors, and are generally found in computing and other compute-intensive applications; they typically form the brains of an entire electrical system. Raw performance is generally the most important metric for microprocessors, though increasingly performance per watt is becoming important as well. Little integration is done in the microprocessor level beyond on-chip cache & I/O, as they almost always sit within a larger system with other ICs to handle supporting functions.

Microcontrollers (MCUs) are smaller standalone processors that perform dedicated or embedded computer functions within a system and usually contain on-chip RAM, ROM, I/O logic and timers in order to minimize cost and to speed design and implementation cycles. Integration is most important, followed by power consumption and performance; though performance will usually vary by application given the diverse use of MCUs.

DSPs are specialized high-speed digital processors typically used to process real-time data. Performance and power consumption are of equal importance. Integration has also been critical in driving DSPs into select applications (like handsets), though many systems use a DSP as a standalone high-speed processor.

Microprocessors

Microprocessors (MPUs) - Digital processors that execute instructions and perform system control functions as programmed with an assembly language. MPUs are optimized for general purpose data processing.
Key Types: x86, PowerPC, ARM, MIPS. 16-, 32-, and 64-bit external bus.

Market Forecast

•Microprocessors are expected to grow slightly from $43.1 billion in 2016 to $45.3 billion in 2021, a 1% CAGR.
•The market is almost entirely PC-based and server based, and is tied to PC unit growth and the mix between desktop, notebooks/Ultrabooks, and servers.

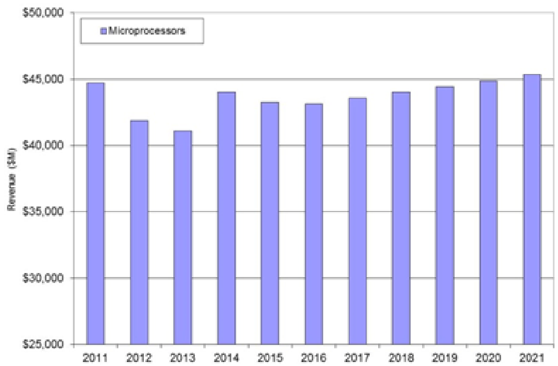

Source: WSTS, IDC, Oppenheimer & Co.

Key Competitors

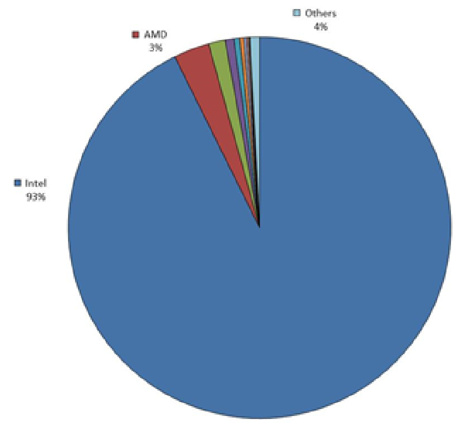

Source: Gartner, Oppenheimer & Co.

PC	Server	Embedded	
Intel	Intel	Intel	Broadcom
AMD	AMD	Cavium	Marvell
		AMD	

A microprocessor (MPU) is a programmable digital processor that executes external instructions and performs system control functions as programmed via software with an assembly language. The architecture is optimized for general purpose data processing and includes an instruction decoder, arithmetic logic unit, registers, cache memory, and additional logic to support operations as determined by the assembly language. An MPU can receive input commands, manipulate data, direct storage of data and initiate application commands to other parts of a system.

The PC and server market represents ~90% of microprocessor revenues. Most PC microprocessors use the x86 CISC (Complex Instruction Set Computer) architecture, which uses microcode to execute very comprehensive instructions. Instructions may be variable in length and use a variety of addressing modes, requiring complex circuitry to decode them. Both Intel and AMD, the major PC microprocessor suppliers, have added their own enhancements to the x86 architecture, both in the actual processor core and the surrounding logic. These enhancements optimize certain common features used in the PC, such as operating system tasks, multimedia functions, and Internet functions.

Microprocessors used for PC peripheral, communications and digital consumer applications—usually referred to as "embedded" processors—generally use either a generic x86 processor or one of many RISC (Reduced Instruction Set Computer) architectures. RISC cores reduce chip complexity by using simpler instructions and keeping the instruction size constant without microcode layer or associated overhead. There are several mainstream RISC architectures, most notably MIPS, PowerPC and ARM. These architectures usually require a license.

Microprocessors are also classified according to bit width, which defines the size of the instructions sent over the external bus (some processors use a higher bit width on the interior of the processor; only the external bus is used for classification). Most PC CPUs use a 32-bit external bus, but can support 64-bit extensions. AMD introduced the 64-bit extension in 2003 and Intel followed suit in 2004; virtually the entire market has transitioned at this point.

In the server space, low-end and mainstream servers have historically used standard 32-bit processors with 64-bit extensions prevalent in the PC market, while high-end servers use proprietary or licensable 64-bit processors (e.g., Sun Microsystems' SPARC or IBM's PowerPC). This has changed in recent years, as Intel's Xeon-based offerings (the vast majority of servers deployed today) are 64-bit systems. Importantly, ARM-based server offerings, from Cavium, and QUALCOMM, have generated market interest as a viable alternative, but have yet to see significant volumes.

In the embedded market, processors use a mix of 8-, 16-, 32-, and 64-bit processors, depending on the application. The highest-end processors would be used in communications infrastructure equipment like softswitches and firewalls, while the lowest-end would be used in consumer electronics and industrial applications. Note that only discrete microprocessors are included in this segment. Most non-PC applications that require a microprocessor actually integrate a licensable RISC core into another ASSP within the system; these are not captured here.

Market Forecast: Microprocessors are heavily tied to the PC market and thus are more dependent on demand for this one application than on broader semiconductor or electronics demand. Overall, we expect microprocessor sales to grow from $44.5 billion in 2014 to about $46.3 billion in 2019, a 1% CAGR.

Competitive Environment: The microprocessor market is one of the most concentrated markets in semiconductors. Intel has effectively maintained a near-monopoly on the PC and server microprocessor market, and therefore held 93% MPU revenue share overall in 2016. Intel's share of the PC and server microprocessor market is 95% (99% market share in Servers), with a somewhat lower unit share in PCs (91%), since its microprocessor ASPs are higher than its competitors'. AMD held 5% revenue share in 2016, with 9% unit share in PCs and 1% servers. AMD competes with Intel in all market segments.

The diagrams below depict revenue and unit share in the PC and server microprocessor markets for 2016.

2016 PC and Server Microprocessor Revenue Share

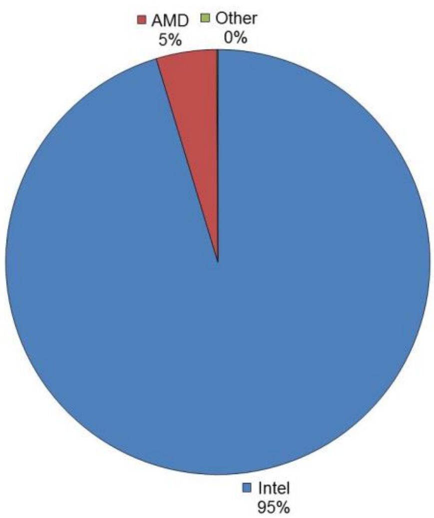

2016 PC and Server Microprocessor Unit Share

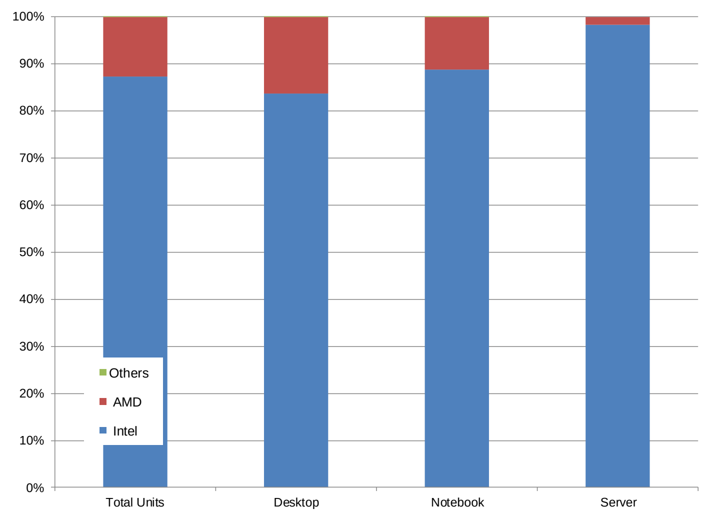

Source: IDC, Oppenheimer & Co.

The embedded microprocessor market is somewhat more fragmented, not surprising given the number of competing architectures and the diversity of applications served. Again, we are counting only discrete MPUs. Intel is the largest player overall with 64% market share within this segment in 2016. Intel's PC processors find many uses in non-PC applications; typically these will be the lower-end or lower-voltage versions of its product portfolio. NXP boasted the No. 2 position with 13% market share within this segment. Within embedded, NXP held the top position in the communications and No. 2 in consumer segments.

Cavium was the No. 3 embedded supplier in 2016, with a 7% share and holds the No. 2 position in communications as the company continues to take share with its multicore product offering. Texas Instruments followed with 5% share, while Renesas secured 3%. Other notable suppliers in this market include AMD, Broadcom, and Marvell.

2016 Embedded Processor Market Share

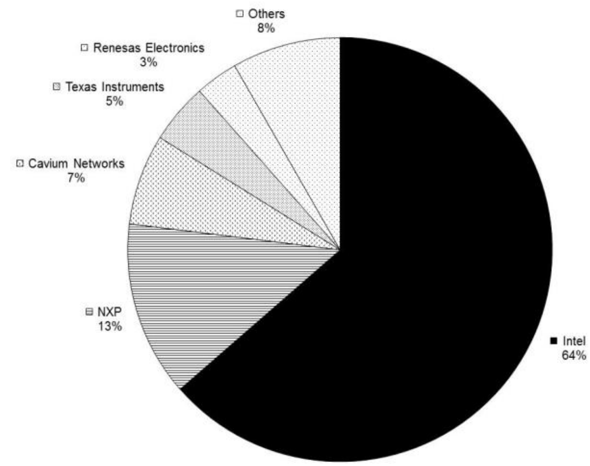

2016 Industrial Embedded MPU Market Share

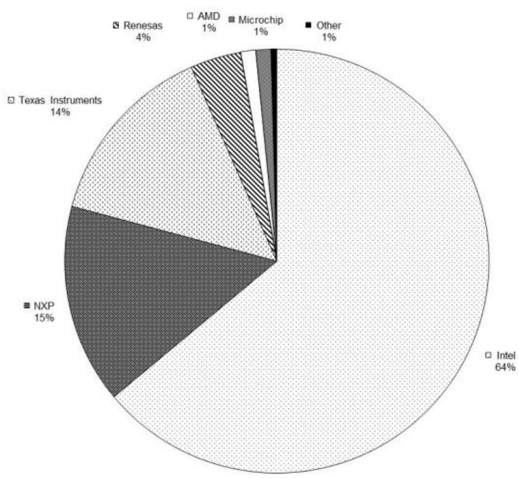

2016 Communications Embedded MPU Market Share

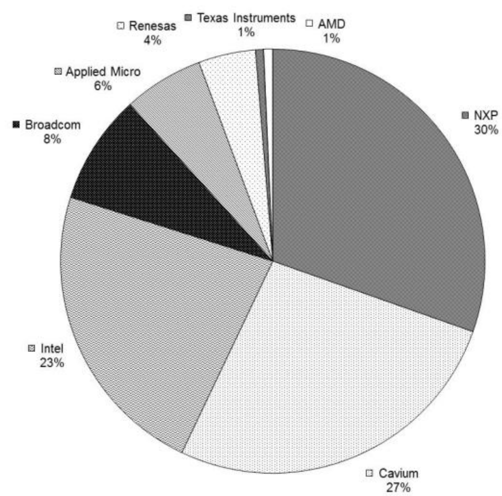

2016 Consumer Embedded MPU Market Share

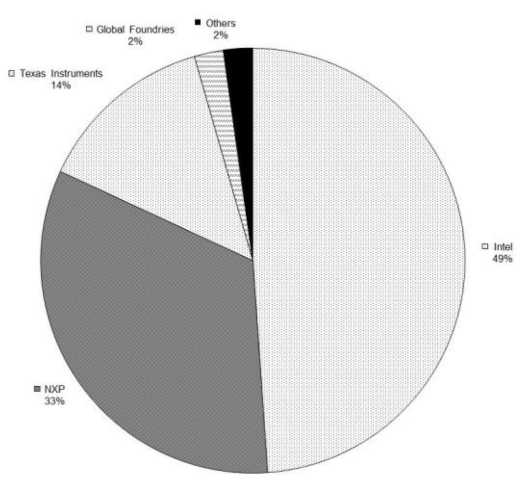

2016 Auto Embedded MPU Market Share

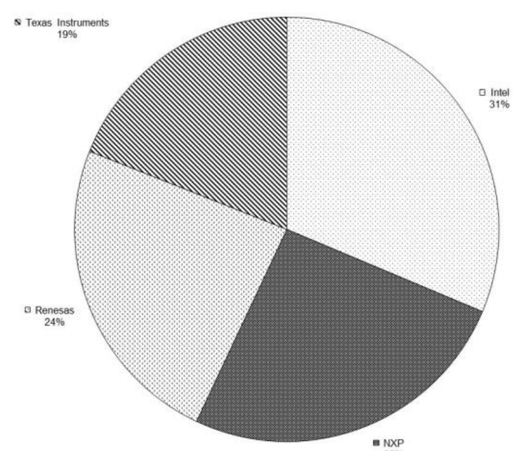

Source: Gartner, Oppenheimer & Co.

Microcontrollers

Microcontrollers (MCUs) – Stand-alone devices that perform dedicated or embedded compute functions within an overall system. MCUs contain single or multiple processing mini-architectures, as well as on-chip RAM, ROM, and I/O logic.
Key Types: 4-, 8-, 16-, 32-bit (or greater). Automotive, consumer, computer, IC card, wireless communication, wired communication.

Market Forecast

- Microcontrollers are expected to grow at a 5% CAGR, from $12.2 billion in 2016 to $15.2 billion in 2021.
- Growth to be led by 32-bit MCUs, which are forecasted to grow at a 7% CAGR, as OEMs look for increased functionality.

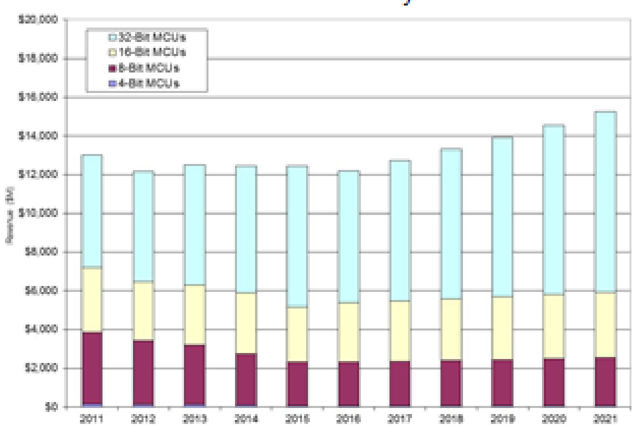

Source: WSTS, Oppenheimer & Co.

Key Competitors

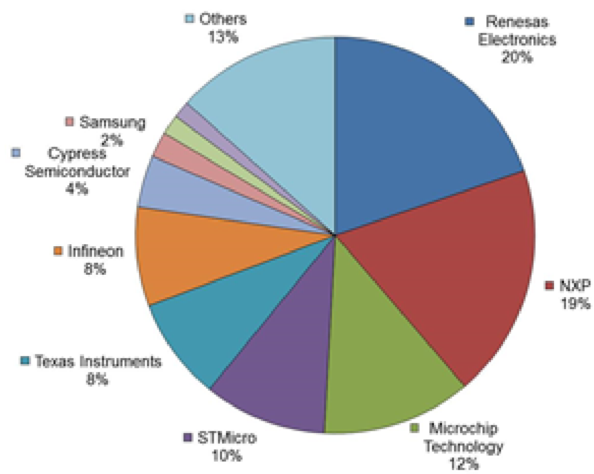

Source: Gartner, Oppenheimer & Co.

Microcontrollers are standalone computation devices that perform dedicated or embedded compute functions within a system. Microcontrollers incorporate simpler processing architectures than microprocessors and are often used to perform only a specific function within a system. They contain some form of RAM and ROM memory as well as additional I/O logic on-chip, and thus can perform relatively simple processing tasks without the need for traditional support ICs found in more complex systems. This lowers the system cost and speeds design and implementation.

Like microprocessors, microcontrollers are classified by external bit width. 4-bit devices are still in use, primarily to control the simplest electronic devices (e.g., a toaster oven) or for very simple embedded functions within a larger system (e.g., raise or lower a car window), but are now less than 1% of the market in terms of revenue. 16-bit has lost share in recent years as the market has bifurcated towards lower-end 8-bit MCUs and higher-functioning 32-bit. Thus 8-bit market share is now 35%, while 16-bit is 22% and 32-bit (or greater) now represents 43% of sales.

By end market, the largest single consumer of microcontrollers is automotive, which accounted for 35% of total microcontroller sales in 2016. Automotive electronics systems perform a diverse array of relatively simple tasks (though higher-end cars have more complex systems), perfectly suited for microcontrollers. Industrial followed and accounted for 24% of the total market. Computing devices were the third largest at 23% of sales.

The exhibits below display microcontroller revenue share by bit width and application.

2016 Microcontroller Revenue Share By Bit Width

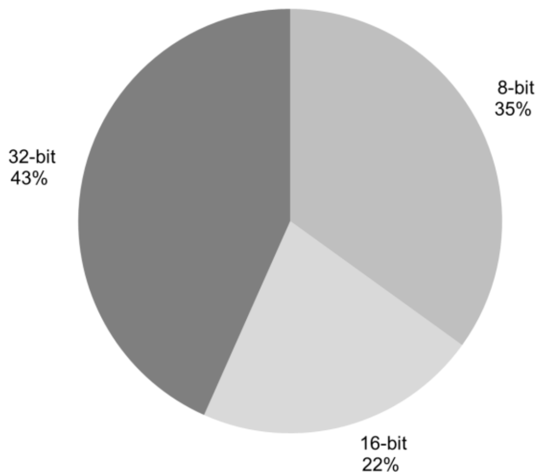

2016 Microcontroller Revenue Share By Application

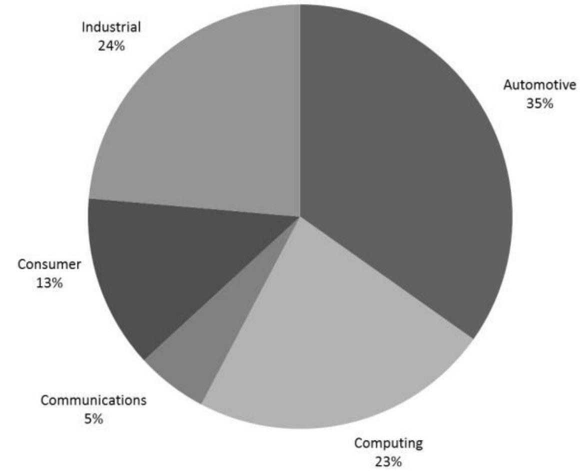

Source: WSTS, Oppenheimer & Co.

Market Forecast: Microcontrollers are expected to see healthy growth given their broad exposure to the electronics landscape. Overall, we expect microcontrollers to grow from $12.2 billion in 2016 to $15.2 billion in 2021, a 5% CAGR. The increasing complexity of automotive and consumer systems should continue to favor the higher bit width segments. 8-bit and 32-bit (or greater) devices are expected to grow at a 4% and 8% CAGR, respectively. 16-bit devices are expected to remain roughly flat over the forecast period.

Competitive Environment: Despite recent consolidation (NXP acquired Freescale, Microchip acquired Atmel, Cypress acquired Spansion), the microcontroller market remains highly fragmented due to the diversity of applications served. The largest supplier in 2016 was Renesas, with 20% market share.

NXP was the No. 2 supplier with a 19% share, and is also the No. 2 supplier in automotive applications. Microchip and STMicro held 12% and 10% of the market in 2016, respectively, while Texas Instruments rounded out the top five with an 8% share. Other major suppliers include Infineon and Cypress. The exhibits below display microcontroller market share, first by bit width and then by application.

2016 8-bit and below MCU Market Share

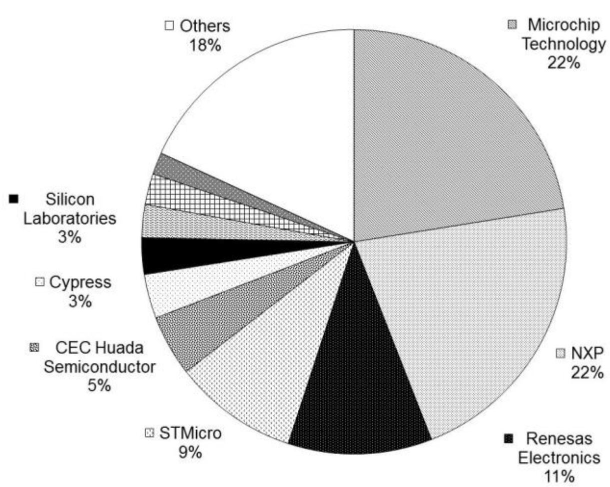

2016 16-bit MCU Market Share

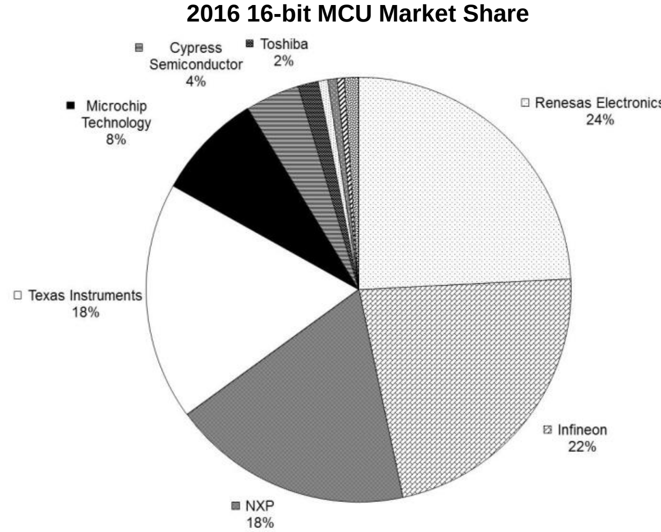

(Exhibit continued on following page)

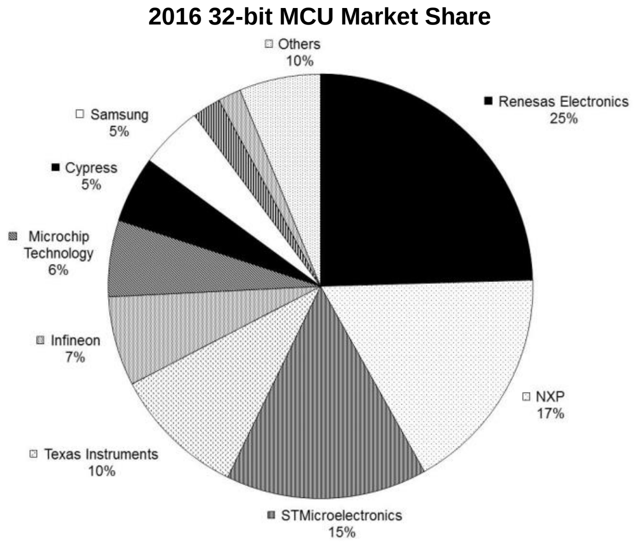

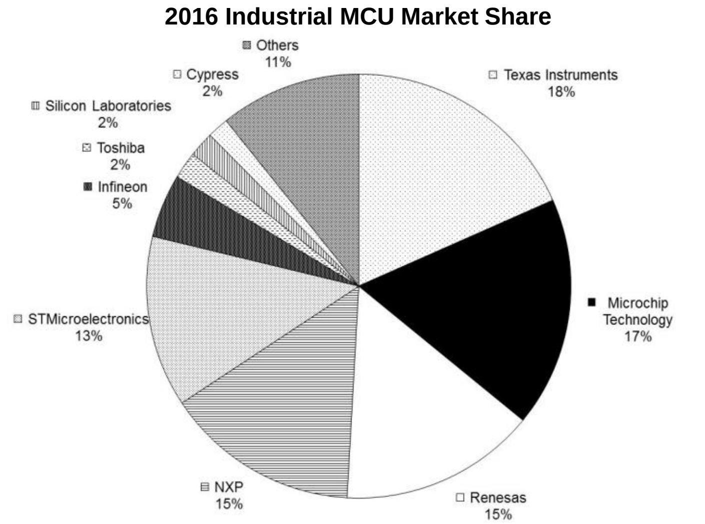

Source: Gartner, Oppenheimer & Co.

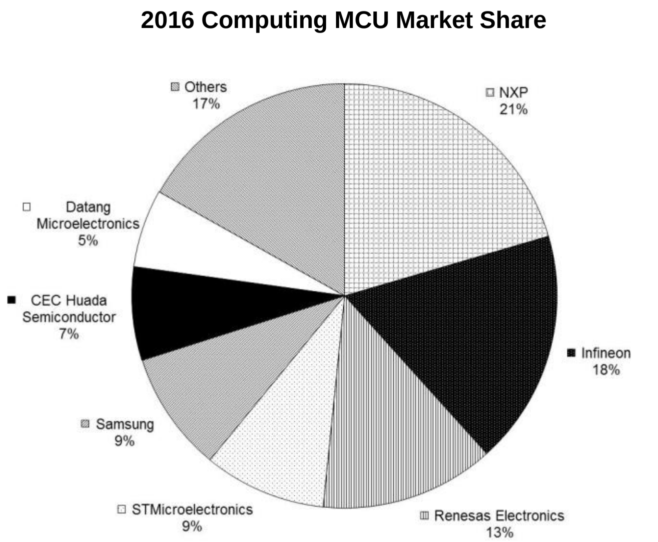

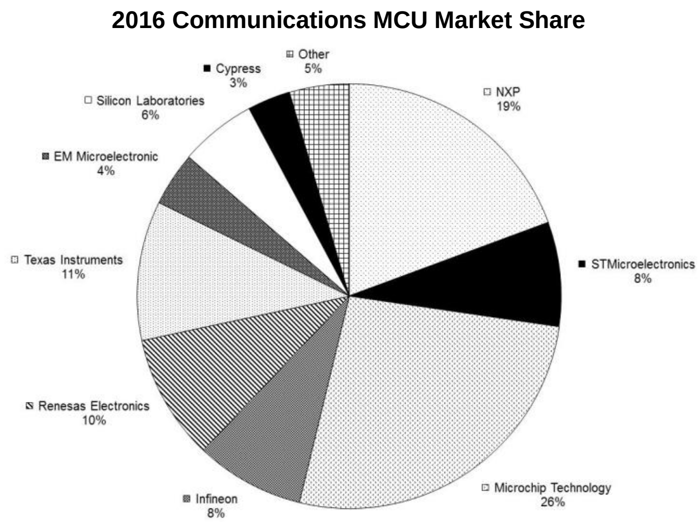

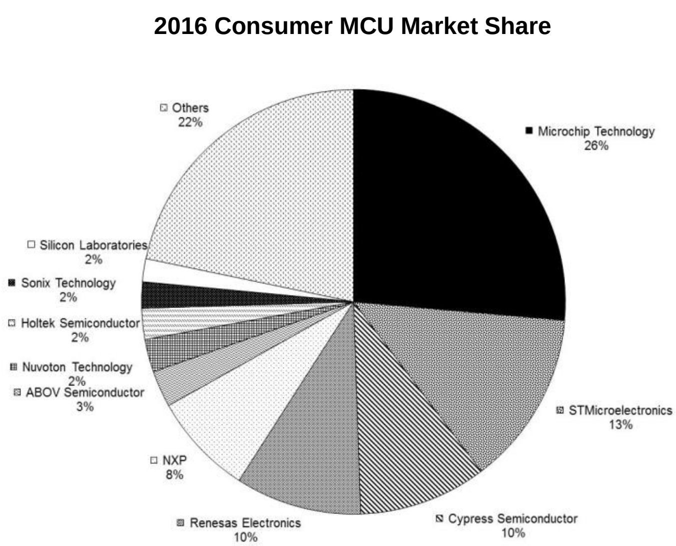

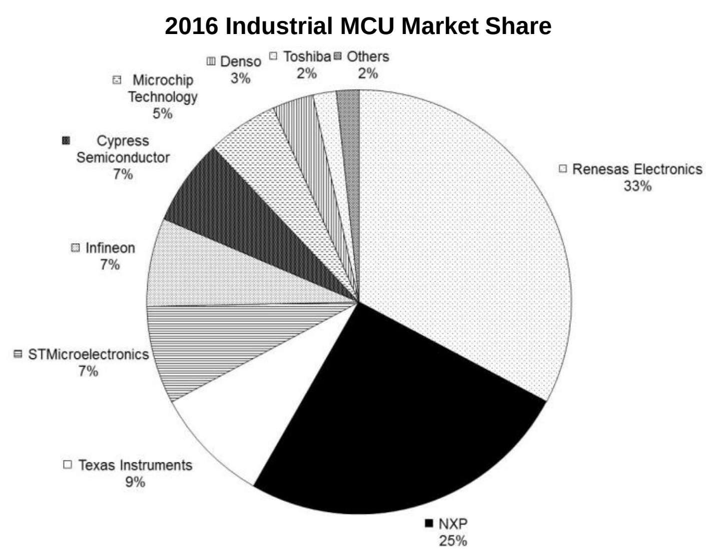

Source: Gartner, Oppenheimer & Co.

Digital Signal Processors

Digital Signal Processors (DSPs) - Specialized high speed programmable processors designed to perform real-time processing of digital signals, such as digital voice and video.

Key Types: Cellular basebands, basestation DSPs, VoIP processors, automotive DSPs.

Market Forecast

•After stabilizing in 2014, DSPs are expected to exhibit modest growth. We expect a 3% CAGR, growing from $ 2.9 billion in 2016 to $3.3 billion in 2021.

•DSPs should trail the overall market as wireless share shifts to vendors of logic based basebands and continued ASP pressure

•Consumer and Automotive remain growth areas.

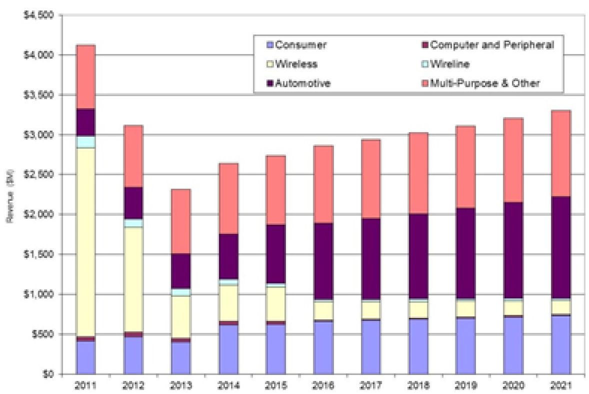

Source: WSTS, Oppenheimer & Co.

Key Competitors

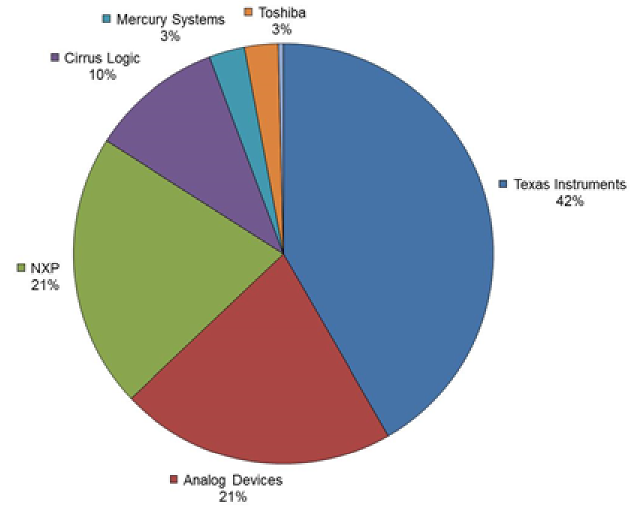

Note: Includes DSPs that are customized for specific applications, but excludes digital logic circuits with embedded DSPs.

Source: Gartner, Oppenheimer & Co.

Digital signal processors (DSPs) are high-speed programmable specialized digital processors that are optimized for performing intensive computations typically required in communications and multimedia applications. DSPs are almost always used in applications that require real-time processing of digital signals. Examples include handset basebands, digital camera image processors, IPTV decoders, and VoIP processors.

By end market, 23% of all DSP revenue was derived from consumer in 2016, the automotive market accounted for 33%, up from just 2% in 2012, while wireless accounted for 8%. Wireline accounted for 1%, mostly broadband modems; 1% sell into the computing market; while multi-purpose grew from 25% of sales in 2012 to 34% in 2016. The exhibit to the right displays 2016 revenue share in DSPs by application.

2016 DSP Segmentation by Application

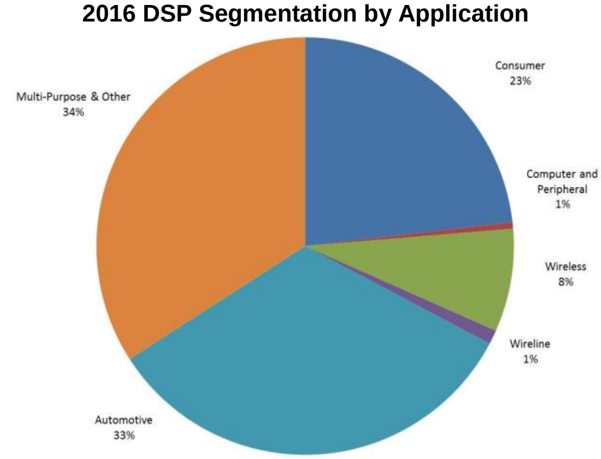

Source: WSTS, Oppenheimer & Co.

Market Forecast: After a rapid decline from 2010 to 2013, DSPs returned to growth in 2014. We expect DSPs grow slightly below the overall semiconductor market growing at a 3% CAGR, from $2.9B in 2016 to $3.3 billion in 2021. Though handset units are growing nicely, a change in competitive dynamics in the wireless market resulted in shifting share to vendors of special purpose logic solutions versus DSPs. TI's technology was classified as DSPs, while the others are classified as special purpose logic. The net of it is we expect wireless DSPs to decline at a 5% rate through 2021.

Automotive is the real growth engine for DSPs, as cars become increasingly more electronic, and OEMs seek to differentiate their offerings through infotainment, safety and convenience features, all enabled by semiconductors. The automotive market is expected to grow at a 6% CAGR through 2021 and increase to 38% of total DSP revenue through 2021 from 33% in 2016.

Competitive Environment: The DSP market is highly concentrated, with four players controlling nearly 95% of revenue in 2016. TI held the No. 1 position in 2016 with a 42% share, with its leadership across the automotive, industrial, consumer and wireless segments. The company completed the wind-down of its baseband business in early 2013. TI has a leading presence in all embedded segments, and is particularly strong in 3G/4G basestations and VoIP equipment.

Analog Devices tied with NXP at No. 2 with 21% share each in 2016. Note that Analog Devices sold its handset business in 2007 to MediaTek, but is still a significant supplier of DSPs for the infrastructure market. NXP (through its acquisition of Freescale) derives much of its leadership position from automotive and wireless base stations.

Note that most special purpose logic devices (discussed later) incorporate some DSP circuits. Most mainstream market segmentations of DSPs exclude these devices, though the lines of distinction are somewhat subjective. For instance Qualcomm makes digital baseband chips for cellular handsets, but they are considered special purpose logic, while TI's had historically been considered DSPs. As mentioned above, the defining line usually rests on how the device is designed. As mentioned above, TI had historically wrapped logic around its standard programmable DSP core; thus, it is classified as a DSP. The same goes for NXP. Qualcomm, on the other hand, designs a baseband ASSP from the ground up and incorporates a DSP-engine; thus, its devices are usually classified as special purpose logic. The same applies to baseband suppliers like Intel and MediaTek. Many other suppliers have exited the baseband business due to intense pricing pressure, leaving only a few major suppliers today.

Digital Logic

Digital Logic:

Devices designed to perform specialized digital processing within a system

Five Key Markets:

Special Purpose Logic

Display Drivers

General Purpose Logic

ASICs

FPGAs/PLDs

Digital logic includes all non-microcomponent logic devices. These devices contain more "hard-coded" logic circuits compared to microcomponents, and thus typically are designed from the ground up for a particular function. Some software programmability is always incorporated, but it is usually to adjust certain details about the device's core functioning (e.g., which features to turn on, how fast to run the device). In general, logic devices trade the benefits of programmability for better performance, lower cost, and better power management.

The largest sub-segment of digital logic is known as special purpose logic, which includes specialized ASSPs designed for a specific application. Only standard "off the shelf" (sometimes referred to as "merchant") devices are included, though some of these ASSPs are customized for a specific OEM customer when there is a large volume opportunity. The other sub-segments within the digital logic category are display drivers, general purpose commodity logic, custom logic (gate array and standard cells, also called ASICs), and field programmable logic devices (PLDs and FPGAs).

Special Purpose Logic

Special Purpose Logic - Application specific digital ICs sold as standard products, designed to service a particular application. Includes ASSPs for PC, display, communication, consumer, and automotive applications.
Key Types: PC core logic, GPUs, HDD SoCs, WLAN/Bluetooth baseband, Ethernet controllers, audio/video decoders, video back end ICs.

Market Forecast

- We expect special purpose logic to grow at a 4% CAGR, from $72.2 billion in 2016 to $88.3 billion in 2021.
- Market growth boosted by integration trends, as digital SoCs incorporate discrete, memory, and analog components, offset by ASP pressures
- Growth led by comms and auto ASSPs.

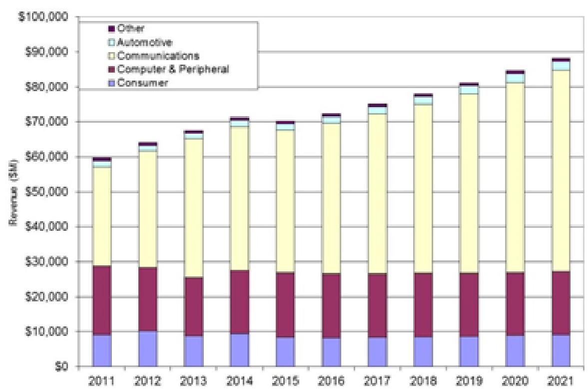

Source: WSTS, Oppenheimer & Co.

Key Competitors

PC and Server
Intel
NVIDIA
AMD

Storage and Peripheral
Marvell
Broadcom
MediaTek
STMicroelectronics
Renesas

Communications
Qualcomm
Broadcom
NXP
Marvell
Cavium

Consumer
STMicroelectronics
NXP
Broadcom
Sony
Toshiba
Renesas
MediaTek

Special purpose logic devices are digital logic circuits sold as standard products, usually designed specifically for a particular application (though not for a specific socket—those are classified as custom ASICs). Designers typically incorporate some degree of system-level knowledge in the architecture, as these devices typically form the core of the equipment into which they are designed.

As discussed in the section on analog ASSPs, many of the highest volume electronics applications are served by a chipset that includes one or more analog ASSPs paired with one or more digital logic devices or ASSPs offered by the same vendor, or by multiple vendors on a single reference design. The analog and digital ASSPs are thus designed in, and subsequently purchased by the OEM, at the same time. A great example of this is a wireless LAN chipset, where most vendors offer a chipset that includes an analog RF chip along with a digital baseband/MAC device. In almost all cases, a design win is awarded for the entire chipset as opposed to individual components. For purposes of this discussion, only the digital device within the chipset is included.

As in the analog ASSP segment, the special purpose logic category is segmented by application; i.e., consumer, communications, computing and peripheral, automotive and multipurpose. In 2016, computing and peripheral logic represented roughly 25% of total market revenue. Communications and consumer logic represented 60% and 12%, respectively, and automotive and other logic totaled 4%. The automotive and other segments are small (but growing rapidly off a small base), as these equipment types are often implemented with microcontrollers as opposed to hard-wired special purpose logic devices.

Market Forecast: The market for special purpose logic is expected to slightly outperform the overall ex-memory semiconductor market as electronic devices use an increasing amount of specialized processing power. There is also a boost from integration, as it is common for digital ASSP vendors to integrate increasing amounts of general purpose logic, embedded memory, and analog circuits into their devices over time. Overall, we expect the market to grow at a 4% CAGR, rising from $72.2 billion in 2016 to $88.3 billion in

2019. Automotive and communications should see the strongest CAGRs at 6% each, while consumer and PCs are expected to be roughly flat.

Competitive Environment: Given the broad set of applications and customers that use special purpose logic, the market is large and rather fragmented. The top suppliers for PC applications include chipset and graphics vendors Intel, NVIDIA, AMD, Broadcom, and Infineon. For storage and peripheral, top vendors include Marvell, Broadcom, STMicroelectronics and Renesas.

Top consumer electronics IC vendors include STMicroelectronics, NXP, Sony, Broadcom, and MediaTek. Note that many of the big vertically integrated Japanese IC makers that supply the consumer electronics market are excluded, since their products would be classified as custom ASICs; the same goes for most suppliers into the video game console and handheld markets.

Top vendors for communications include wireless IC leaders Qualcomm and MediaTek. Marvell had been a player in the space, but de-emphasized wireless ICs in recent years, following Broadcom's lead from 2013. On the wireline side, key ASSP vendors include Broadcom, Texas Instruments, Intel, Marvell, and Qualcomm (CSR).

The table below provides additional detail, listing some high-volume special purpose logic devices and their vendors.

Market Segment	**Device**	**Key Vendors**
Computing & Peripheral	PC core logic	Intel, NVIDIA, AMD, Broadcom
	Graphics controllers	NVIDIA, AMD
	HDD SoCs & hard disk controllers	Marvell, Broadcom, STMicroelectronics, Infineon, Renesas
	Printer SoC	STMicroelectronics, Texas Instruments, Marvell, Microsemi, Conexant, Broadcom, MediaTek
Communications	Cellular baseband	Qualcomm, Intel, MediaTek, Samsung
	Ethernet controllers and switch ICs	Broadcom, Marvell, Intel, Realtek, Microsemi
	Wireless LAN baseband/MAC	Broadcom, Qualcomm, Intel, Marvell, MediaTek, Realtek, Samsung
	Bluetooth baseband/MAC/SoC	Broadcom, Qualcomm, Intel, Marvell, MediaTek, Realtek, Samsung
	Modem CPE SoC	Broadcom, Infineon, Intel, Qualcomm
	Telecom/Datacom	Microsemi, Broadcom, Infineon, MA-COM, Inphi, Mellanox
Consumer	Set-top box back-end	STMicroelectronics, Broadcom, NXP, Maxlinear
	Digital TV back end	NXP, Maxlinear, Sharp, MediaTek, Renesas, Broadcom
	Digital camera image processor	TI, MediaTek, Ambarella, Renesas

Display Drivers

<u>Display Drivers</u> - Standard logic devices designed to control and drive flat panel displays such as LCD or plasma.
Key Types: LCD driver, plasma display driver, source driver, gate driver.

Market Forecast

- We expect display drivers to grow at a 3% CAGR, from $6.0 billion in 2016 to $7.0 billion in 2021.
- Slowing smartphone growth and ASP pressures drive below-market growth.
- Flat panel TV, notebook monitors, smartphones and tablets are the primary growth drivers.

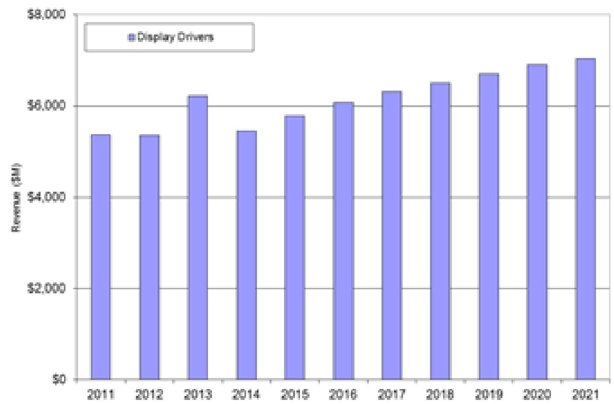

Source: WSTS, Oppenheimer & Co.

Key Competitors

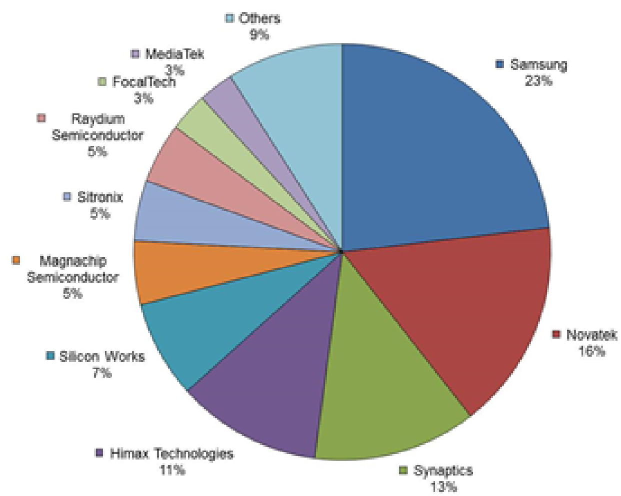

Source: Gartner, Oppenheimer & Co.

Display drivers are specialized logic devices designed to control and drive flat panel displays such as LCD, plasma or OLED. Traditional CRT display technology uses a phosphorous coated screen to trap light flashed from a light source; flat panel displays, in contrast, use semiconductor driver ICs to manipulate pixels on an LCD or plasma screen.

There are two basic kinds of drivers: source and gate. Source drivers sit across the top of the panel and generate signals that penetrate individual columns of pixels. Three source drivers are needed for each column of pixels: red, green and blue. Gate drivers run along the side of the panel and manipulate individual rows of pixels. Signals generated by the gate drivers intersect signals driven by the source drivers, turning individual pixels on or off.

Market Forecast: We expect display drivers to grow at a 3% CAGR, from $6.0 billion in 2016 to $7.0 billion in 2021. LCD screens are an integral part of a category that itself is growing nicely, such as smartphones, tablets and portable gaming. Decelerating smartphone unit growth likely drives slower growth than in the past five or ten years. We view Ultrabooks/Notebooks and TVs as a likely flattish unit opportunity at best going forward. In all cases, price pressure is a major factor weighing on the revenue CAGR.

Competitive Environment: Display drivers are quasi-commodity parts, though driver performance does differentiate products from one vendor to the next. The top suppliers in 2016 included Samsung, with a 23% market share, Novatek with 16%, Synaptics (through its Renesas DDIC business acquisition) with 13%, Himax with 11% and Silicon Works with 7%. Other significant suppliers include Magnachip, Sitronix, Raydium, FocalTech and MediaTek.

General Purpose Logic

General Purpose Logic - Standard commodity catalog logic parts designed for use in multiple applications in various end markets.
Key Types: Simple gates, flip-flop circuits, switches, registers.

Market Forecast

•We expect general purpose logic to grow at a 2% CAGR over the next few years, rising from $1.5 billion in 2016 to $1.7 billion in 2021.
•Growth should trail overall semi market growth as integrated ASSPs eliminate the need for separate logic circuits in many applications.

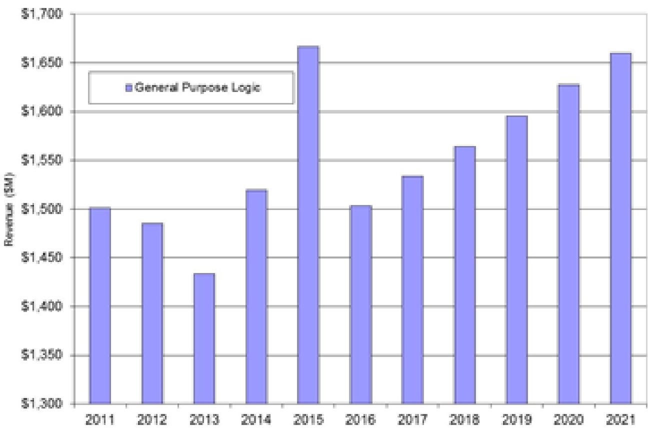

Source: WSTS, Oppenheimer & Co.

Key Competitors

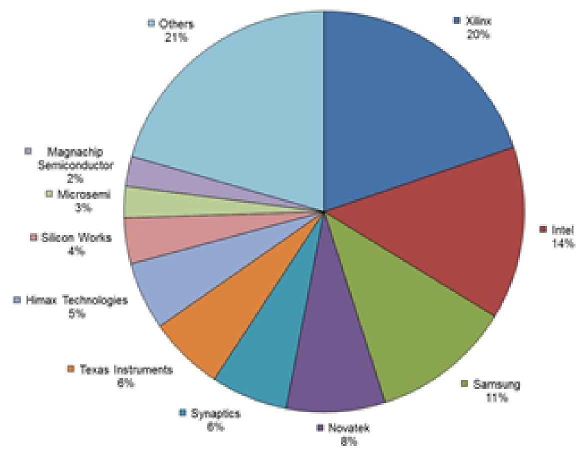

Source: Gartner, Oppenheimer & Co.

General purpose logic includes standard logic parts designed for use in multiple applications in various end markets. Though designed to perform a specific function, these ICs are not tailored for use in any one application. Parts in this category are usually very low ASP, very high-volume devices; examples include simple gates, flip-flop circuits and registers; they are typically used to assist the functioning of larger logic or processing devices. Many of these parts are commodities, though some are more proprietary.

Market Forecast: The market for general purpose logic is heavily tied to the overall semiconductor cycle, as these devices sell into every end market and application in the electronics landscape. They are also heavily commoditized and pricing can compress severely in down-cycles and rise in up-cycles. We expect growth in this segment to track in line to modestly below the overall ex-memory semiconductor sector, as ASSP vendors continue to integrate additional functionality into their devices that eliminate the need for external general purpose circuits in many applications. Overall, we expect a 2% CAGR through 2021, with the market growing from $1.5 billion in 2016 to $1.7 billion in 2021.

Competitive Environment: The general purpose logic category includes a wide variety of device types, but most are commodities, and thus a few dedicated players dominate the market. Most of these are vertically integrated IDMs that use standard logic as a fab filler, though a few niche players are fabless. The top standard logic producer in 2016 was Xilinx with a 20% share, Intel (through its Altera acquisition) with 14%, Samsung with 11% and Novatek with 8%. Other key players include Himax, TI (which has de-emphasized this business), Synaptics, and Silicon Works.

ASICs

ASICs (Application Specific Integrated Circuits) - Customized ICs designed for specific customers to be used in specific platforms.
Key Types: Standard cell, gate array.

Market Forecast

- We expect the ASIC market to grow at a 4% CAGR, from $12.2 billion in 2016 to $14.7 billion in 2021.
- Game consoles are a volatile driver; in most other applications the trend away from ASICs should continue.
- The shift from gate arrays to standard cells should continue as OEMs demand higher performance devices.

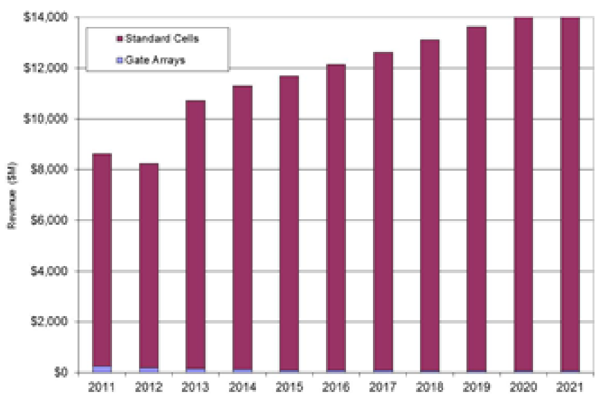

Source: WSTS, Oppenheimer & Co.

Key Competitors

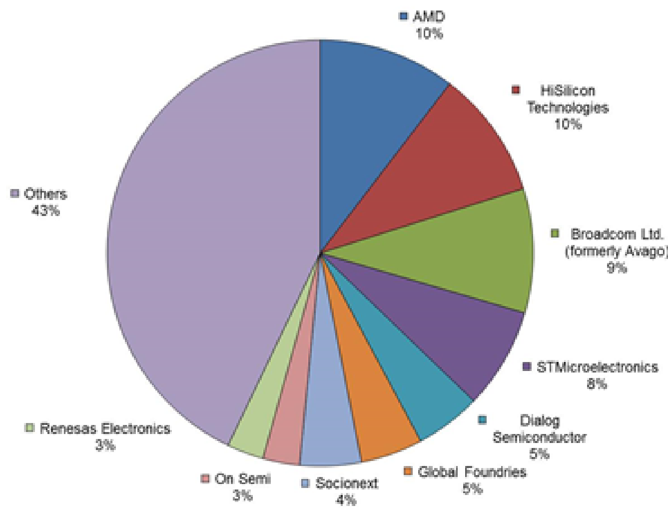

Source: Gartner, Oppenheimer & Co.

Moving away from standard logic devices and into the custom-designed realm, the ASIC (application specific integrated circuit) market includes all customized integrated circuits designed for specific customers to be used in a specific platform. The key distinction with an ASIC is that the OEM plays a heavy role in designing the device. The most prominent examples of ASICs are in the game console market, where most silicon is designed for a single platform, which, given the high volume it is likely to generate, justifies the cost of designing a custom chip. Note that although classified in the logic category, ASICs can include analog, digital, or memory circuits, as long as they are custom designed. Most ASICs, however, are pure logic devices anyway.

ASIC devices can either be designed by the equipment manufacturer internally—many OEMs maintain sizable chip design teams—or contracted out to an ASIC house. The level of cooperation between an OEM and an ASIC house can vary from program to program.

ASICs are divided into two broad categories:

- **Gate Arrays**: Circuits consisting of fixed and regular arrangements of transistor cells forming a matrix of logic gates of various standard densities. An OEM using a gate array uses a series of software tools and macro libraries to design the mask for the top layers of the device, which determines its functionality.

- **Standard Cells**: Also known as “cell-based ICs,” these circuits consist of user-specified arrangements of predefined and fixed sub-circuits of any function (analog, logic, memory). These ICs are designed from the ground up and use standard dimensions of components or gates to pack them together more uniformly. These are more costly than gate arrays but are more efficient.

ASICs tend to be the first solutions to be used in emerging electronic equipment markets, as most electronic devices require some level of specialized digital or analog processing, but few standard, off-the-shelf components exist during the introductory phase of the technology life cycle. These ASICs tend to be costly to design, but usually deliver the performance necessary to power the device given the custom design; and the costs can be spread over the lifetime of the product if it takes off. Also, given the early stage of the technology life cycle, the OEM is likely earning high margins, justifying the cost of the ASIC.

As the technology matures and the market grows, other OEMs will likely enter the market with competitive offerings that either improve upon existing platforms or offer the same functionality at lower cost, or both (usually both). These competitive offerings will sometimes use newer ASIC solutions. Other times, they will leverage standard products (ASSPs) that may have come to market a few years into the life cycle of the emerging technology by IC vendors who were attracted by the growth prospects of the market as it took off. Over time, more and more OEMs will migrate to ASSP solutions, and the ASICs will disappear (and the ASIC vendors will go on to the next new market).

Although there will always be a need for ASICs, most OEMs will choose a standard product if it is available. This is the primary reason that the transition from ASICs to ASSPs is a constant theme in semiconductors (though it varies from market to market based on the stage of maturity). There are two notable exceptions: the first is where a single platform is likely to do so much volume that it justifies having a custom part designed for it. A prime example of this is the game console application mentioned above. The second case is where silicon performance is such a differentiating factor for the OEM that, to retain its competitive advantage, OEMs will develop ASICs instead of leveraging outside IC designers. A good example is Cisco, which views the performance of its switches and routers to be so integral to its strategy that it employs several hundred ASIC engineers.

Market Forecast: Overall, we expect the market for ASICs to grow at a 4% rate through 2021, rising from $12.2 billion in 2016 to $14.7 billion in 2021. We expect the ASIC CAGR to perform roughly in line with the overall semi market, given that OEMs are increasingly looking to differentiate their offerings and integrate multiple systems onto one die, while also looking for lower-cost alternatives. Historically growth in the ASIC market has been buoyed by demand from gaming consoles; however, this is no longer the case.

Competitive Environment: As described above, there are a number of types of relationships in the ASIC market. Some OEMs are vertically integrated, and design and manufacture their own ASICs completely internally (this is common in Japan). On the other end of the spectrum, some OEMs contract with ASIC houses and leave the chip design and manufacture to the partner. Others fall somewhere in the middle, with the OEM maintaining an internal semiconductor design team, but partnering with an ASIC vendor for some of the IP as well as for gate construction and fabrication. This yields a vendor list that includes both OEMs (Samsung, IBM, Toshiba, Fujitsu), IDMs (Renesas, Texas Instruments), and even fabless chip makers (Avago/LSI).

The market landscape has become quite diverse. The largest ASIC houses in 2016 were AMD and HiSilicon with just 10% share each. AMD’s share is boosted by its gaming console business and HiSilicon is a fully owned subsidiary of Huawei. Other key vendors include Broadcom, STMicro, Dialog, and Global Foundries.

FPGAs/PLDs

Programmable Logic Devices (PLDs) – Devices that contain arrays of programmabl logic gates that can be configured into higher level logic devices by the customer during system integration.
Key Types: Field Programmable Gate Array (FPGA), Programmable Array Logic (PAL), Complex Programmable Logic Device (CPLD).

Market Forecast

•We expect FPGAs to grow at an 5% CAGR, rising from $4.7 billion in 2016 to $6.0 billion in 2021.
•While communications will remain a major driver, the emergence of low-cost, high-density FPGAs is broadening the application base for FPGAs.

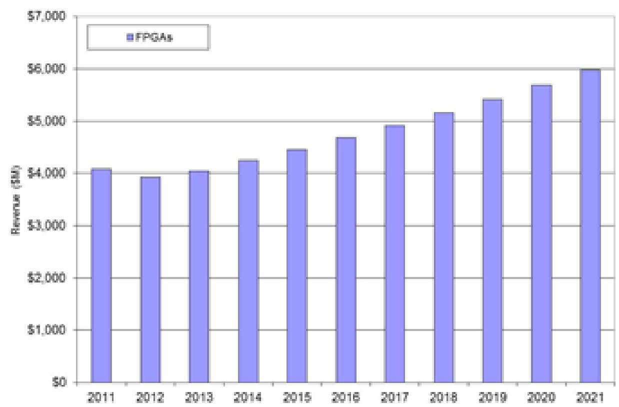

Source: WSTS, Oppenheimer & Co.

Key Competitors

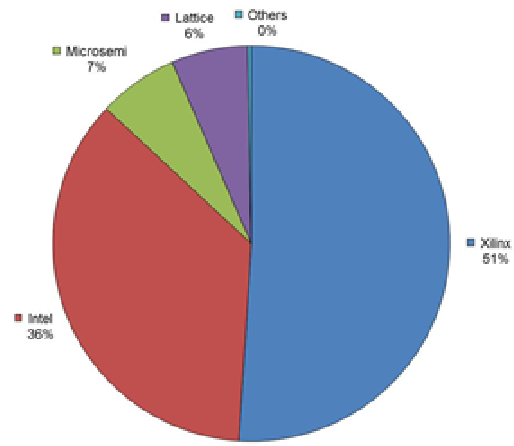

Source: Gartner, Oppenheimer & Co.

Programmable logic devices (PLDs) are devices consisting of arrays of programmable logic gates that can be configured into higher level logic patterns by the customer during system integration. With a PLD, the OEM purchases a catalog part and uses the PLD vendor's software to design and configure the top layers of the device to perform the desired function. The OEM then inputs the architecture into the device, which stores it and adopts it for permanent use (some PLDs can be programmed multiple times). PLDs therefore provide the customizability of an ASIC without the need to design and fab new devices for each platform. There are several key types of PLD devices:

- **Field Programmable Gate Array (FPGA)**: High-density PLD containing small logic cells interconnected through a distributed array of programmable switches. This type of architecture produces statistically varying results in performance and functional capacity, but offers high register counts.
- **Programmable Array Logic (PAL)**: Simple PLD devices with a small number of gates and a set number of inputs and outputs. Often used for glue logic.
- **Complex Programmable Logic Device (CPLD)**: High-density PLD consisting of macrocells that are interconnected through a central Global Routing Pool. This type of architecture provides high speeds and predictable performance, and is generally the preferred architecture for implementing high-speed logic.

In terms of end market use, PLDs are closely tied to the communications and industrial markets, which represented 41% and 25% of 2016 PLD revenue, respectively. Consumer applications have not historically been large users of PLDs, though advanced set-top boxes are adopting them, and in 2016 consumer represented 8% of the PLD market. Military/aerospace and data processing, primarily storage and server applications each consumed an additional 14% and 6% of 2016 revenue, respectively. Automotive represented just 5% of the market. The exhibits below display the end market usage for PLDs as well as top PLD applications by end market.

2016 PLD Revenue Share By Application

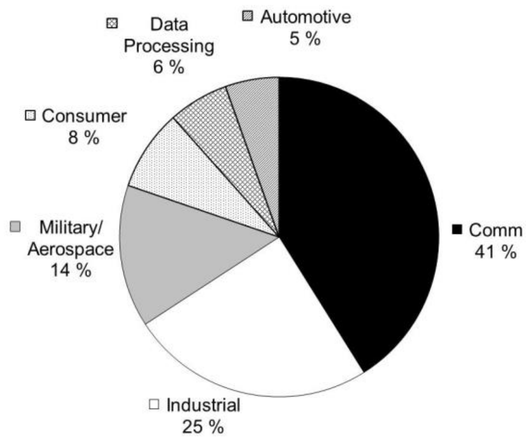

Top PLD Applications, By End Market

Data Processing	Wireline	Consumer	Industrial
Servers	Routers	LCD TV	Manufacturing systems
Enterprise storage	Broadband infrastructure	Digital CRT TV	Instruments
Printers	OTN	DVRs	Medical equipment
	SONET add/drop multiplexers	Set-top boxes	Security/ energy management
Wireless	MSPPs	Digital cameras/ camcorders	
Infrastructure (baseband subsystem)	LAN switches		
	SAN switches		
	CO switches		
	PBXs		
	VoIP gateways		

Source: Gartner, Oppenheimer & Co.

Market Forecast: Programmable logic has become an important solution in advanced wireline and wireless communications, with this end market representing the largest chunk of PLD revenue (~41% of PLD revenues in 2016). As a result, demand for PLDs was very volatile over the last two decades, hit especially hard in 2001 and 2002 when demand for components in the wired telecom infrastructure market dropped. The market rebounded strongly in 2003, and subsequent years have been characterized by more consistent demand as wireline and wireless infrastructure demand stabilized.

Looking ahead, PLD vendors remain focused on the communications market, but are also broadening their product offerings to include smaller, lower-cost PLDs to compete in consumer, industrial, data processing applications that have traditionally favored ASICs or ASSPs. Artificial intelligence is becoming a vertical that PLD/FPGA manufacturers are focusing on as the chips can be used to accelerate training or inference calculations. While communications should continue to outweigh this opportunity, these new markets should add to the long-term SAM for PLDs. In total, we expect programmable logic revenues to grow at a 5% rate, rising from $4.7 billion in 2016 to $7.0 billion in 2021.

Competitive Environment: The PLD market is highly concentrated and has undergone a fair amount of consolidation in the past few years as the complexity and size (in terms of number of transistors) have steadily increased. In 2016, the top two vendors held ~87% share. Xilinx was the No. 1 vendor in 2016 with 51% share. Intel, through its acquisition of Altera, followed at No. 2 with 36% share and Microsemi held 7%. Lattice rounded out the landscape with a 6% share.

Memory

Memory:

Devices designed to store electrical data, either temporarily or permanently

Five Key Markets:

Volatile Memory	Non-Volatile Memory
DRAM	NOR Flash
SRAM	NAND Flash
	Legacy Non-Volatile Memory

Memory devices are digital semiconductor devices designed to store data in an electronic format on either a temporary or permanent basis. Though they are digital in nature, the manufacturing processes used to build memories are often quite different from a bulk CMOS process, especially in the case of DRAM and NAND flash. In nearly all cases, they are optimized for cost and high-volume production.

Volatile memory includes those memory devices that lose their stored information when they lose power. The most important types are DRAM (dynamic random access memory) and SRAM (static random access memory).

Non-volatile memory retains the stored information when the system is powered off. The most important types are NOR flash, NAND flash, mask ROM, EPROM and EEPROM (we group the latter three in a “legacy” category).

Note that semiconductor-based storage is often referred to as “solid state” storage, in contrast to magnetic (hard disk drive), optical (CD and DVD), or tape storage. Semiconductor-based storage is more expensive on a per-bit basis than all of these media, but it is faster and far more flexible in terms of its use.

DRAM

DRAM (Dynamic Random Access Memory) - The most common type of volatile memory, used primarily as system memory in PCs and servers.
Key Types: SDRAM (Synchronous DRAM), DDR (Double Data Rate), DDR2, DDR3, DDR4, QDR (Quad Data Rate), RDRAM (Rambus DRAM).

Market Forecast

- We expect DRAM revenue to increase at a 7% CAGR, from $41.2 billion in 2016 to $57.7 billion in 2021, following significant growth on favorable pricing conditions in 2017
- The market should remain volatile, up the most in boom years and down the most in downturns; pricing has historically swung wildly based on supply/demand trends.

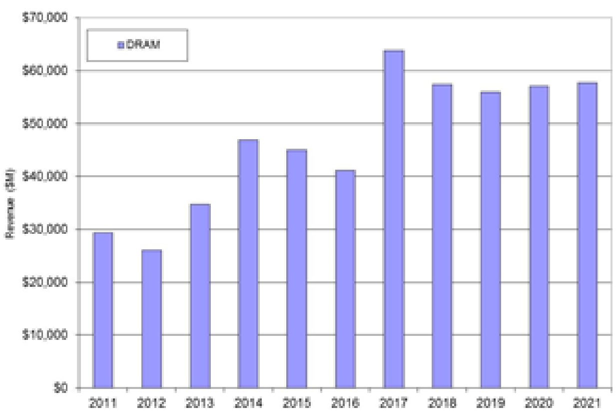

Source: WSTS, Oppenheimer & Co.

Key Competitors

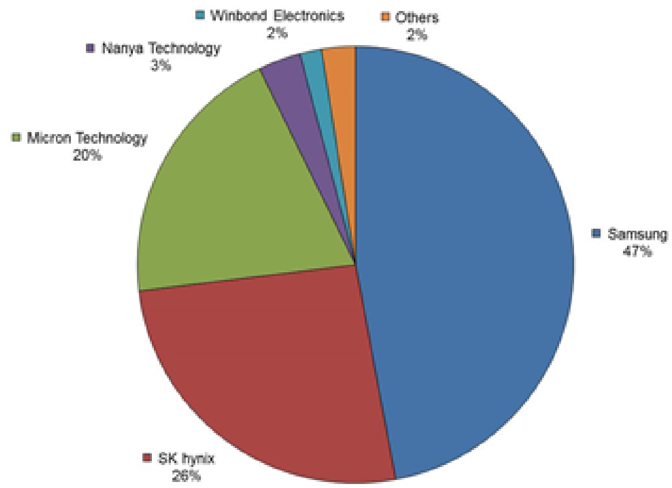

Source: Gartner, Oppenheimer & Co.

DRAM (dynamic random access memory) is the most common type of volatile memory; it is digital memory that loses its stored information once the power is removed. The term "dynamic" means that DRAM, unlike SRAM, must be continually refreshed by the system, slowing down memory access but allowing for a low-cost design and easy scalability to large densities. Like most memory, bit words in DRAM can be written, stored and read in any sequence.

DRAM is used in all electronics markets, but has an especially high concentration in computing and communications applications, which also tend to use the fastest and highest density DRAM devices. In 2016, 47% of DRAM revenue was derived from PCs and servers. Communications represented 39% of the market, while Industrial, auto, and consumer applications each represented less than 10% of DRAM revenue, as these applications consume relatively small amounts of DRAM or are better suited to SRAM or NAND.

The exhibits below display 2016 DRAM unit shipments by application

2016 DRAM Revenue By Application

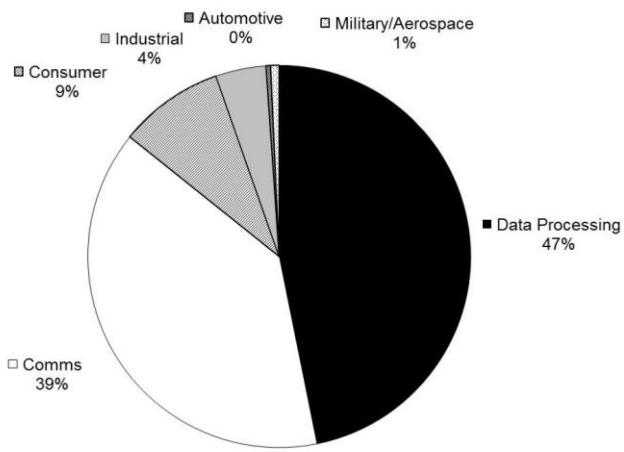

Source: Gartner, IDC, Oppenheimer & Co.

In addition to being classified by interface type, DRAM is quoted by density or bit capacity, with the most common size in use today being 4 Gb (63% of units in 2016). 8 Gb represented 20% of units in 2016, 2 Gb Mb devices represented 8%, 1 Gb was 6% while 512mb still represented ~2% of units, though they are being phased out of most applications given rapid changes in storage requirements.

The exhibits below display 2016 DRAM unit shipments by density, as well as historical and forecast unit shipment data showing the steady migration to larger densities over time.

2016 DRAM Unit Shipments By Density

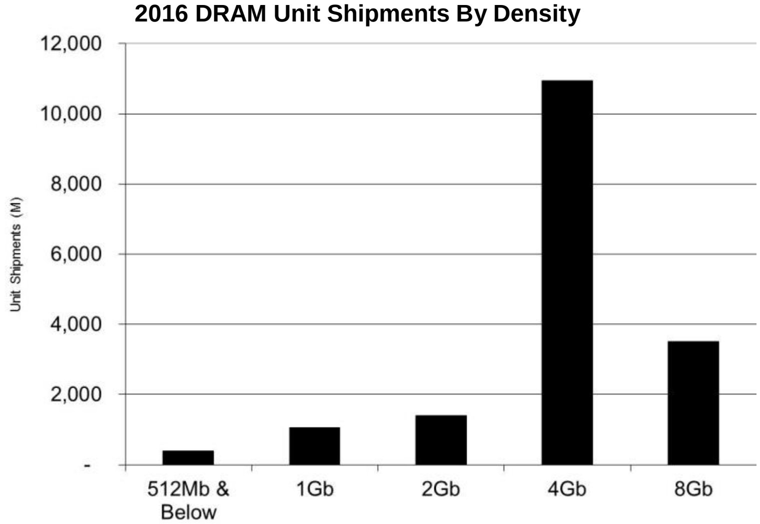

DRAM Unit Shipments By Density, 1997-2019

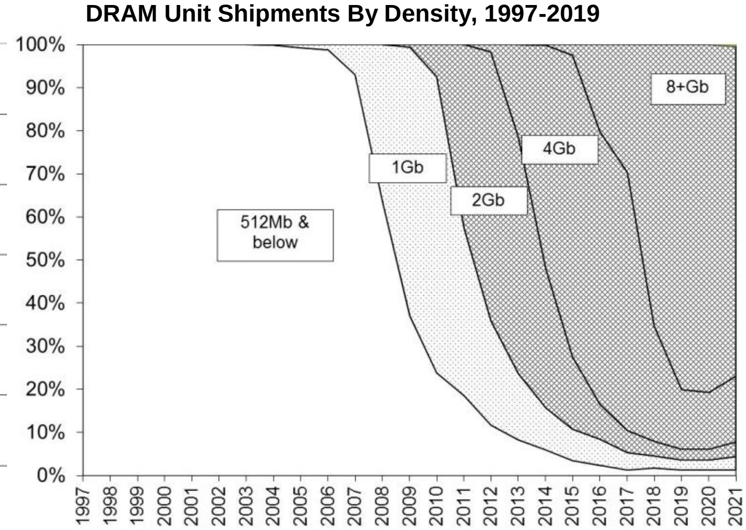

Source: Gartner, Oppenheimer & Co.

Note that most DRAM sold into PCs and servers is not sold in chip form, and instead is soldered onto memory modules. A standard SIMM (Single In-line Memory Module) or DIMM (Dual In-line Memory Module) module is a small PCB (printed circuit board) containing eight or 16 DRAM chips (there are eight bits to a byte, so, for example, a 256 MB DRAM module would usually have eight 256 Mb DRAM chips or 16 128 Mb DRAM chips). A DRAM module is pictured on the next page.

In the server market, memory must be more reliable and robust, and therefore generally includes an extra DRAM IC for buffering along with a timing device to perform register functions; these are known as registered DIMM (R-DIMM) modules. In 2006, Intel-based servers began migrating to fully buffered DIMMs (FB-DIMMs), which replace the register with a high-speed serial I/O interface IC known as an AMB (Advanced Memory Buffer) directly on the module for interface to the server chipset.

PC DRAM Module

Source: TechPowerUp, Oppenheimer & Co.

Market Forecast: DRAM is a commodity product, with most ICs conforming to specific industry-standard densities, interfaces and power metrics. There are high fixed costs in its production, and most vendors make investments in manufacturing well in advance of actual demand, and run their factories close to full. As such, the market is subject to periods of steep price inflation during expansions as capacity tightens, and precipitous price declines during downturns as vendors drop prices to levels that approach variable cost.

In 2001, the DRAM market saw its worst year ever, as DRAM demand dropped on weak PC sales against a backdrop of significant oversupply; prices cratered as vendors continued to fill their fabs despite slower demand. The market recovered in 2003-04 as PC demand rebounded and as OEMs took advantage of the low cost of DRAM by aggressively ramping the number of megabytes per PC. The market worsened slightly in 2005 despite relatively healthy PC demand, as pricing softened due to moderate oversupply. 2006 was a very strong year for DRAM, up 32%, as memory pricing increased in anticipation of the launch of Vista. 2008 was especially difficult (-23% y/y) following a strong consumer pullback in spending. After the 2008-2009 downturn, capacity tightened and demand snapped back, leading to DRAM increasing an impressive 75% year-over-year in 2010. The boom was short lived and DRAM declined significantly in both 2011 (-25% y/y) and 2012 (-11% y/y) as PC units (and DRAM pricing) came under pressure from smartphones and tablets.

The DRAM (and memory market in general) saw massive consolidation in 2013 and 2014, thus pricing stabilized, and DRAM revenue grew 33% in 2013 and another 34% in 2014. After declining 4% and 8% in 2015 and 2016, respectively, favorable pricing conditions driven by strong demand and supply constraints led to significant growth in 2017. 1H17 volatile memory was up nearly 75% Y/Y.

Looking ahead, the DRAM market should remain volatile, up the most in boom years and down the most in downturns; pricing should be volatile, varying with supply trends. Overall, we expect the market to increase modestly from $41.2 billion in 2016 to $57.7 billion in 2021, a CAGR of 7% following 2017's significant growth.

Competitive Environment: With the exception of Micron in the United States, all major DRAM vendors are located in Asia. The top vendor in 2016 was Samsung, which held an impressive 47% share. Hynix held 26% share and Micron (through its Elpida acquisition), was No. 3 with 20%. Rounding out the top vendors were Taiwanese DRAM players Nanya (3% share) and Windbond (2% share).

NAND Flash

NAND Flash – Type of non-volatile rewritable memory that has very fast read and write times, but is accessed in blocks and so cannot support execute-in-place. Used primarily for bulk data storage.
Key Types: SLC, MLC, TLC.

Market Forecast

- We expect the NAND flash market to grow at a 9% CAGR, from $32.0 billion in 2016 to $48.9 billion in 2021.
- NAND is benefiting from the growth of digital consumer devices and is also encroaching on NOR in handsets.
- NAND and DRAM cycles are related as fabs are easily converted from one to the other as demand fluctuates.

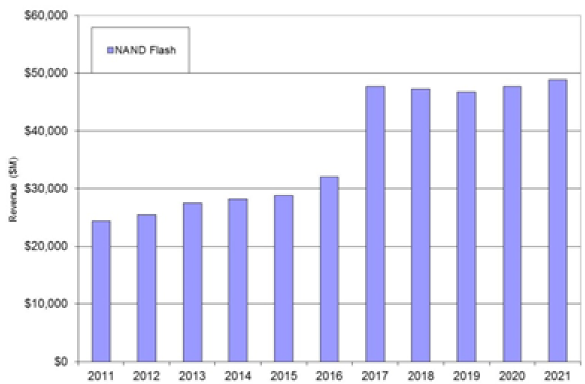

Source: WSTS, Oppenheimer & Co.

Key Competitors

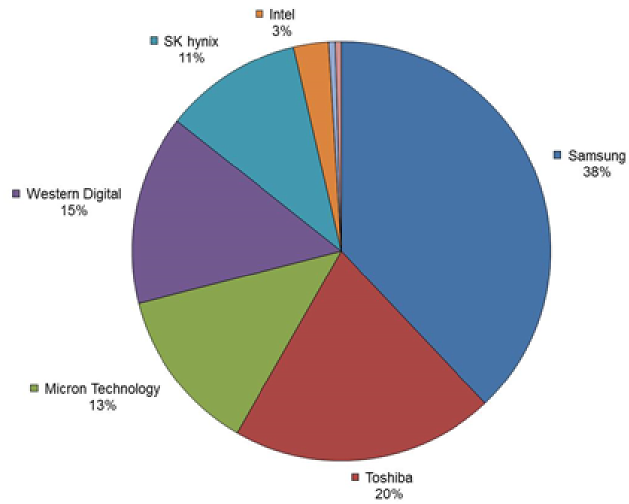

Source: Gartner, Oppenheimer & Co.

The largest type of flash memory is known as NAND flash (about 95% of total flash revenue in 2016). Although they were both invented by Toshiba and share the name "flash," NAND is quite different from NOR in its construction and operation. NAND cells use tunnel injection for writing and tunnel release for erasing, and are organized in blocks of 512 or 2,048 bytes. The blocks are accessed in series similar to hard disk drives, and therefore cannot support random access or execute-in-place.

NAND is far cheaper to produce than NOR and scales more easily to larger densities. The manufacturing process is actually somewhat similar to a DRAM process, and many NAND flash fabs are actually converted DRAM fabs, making their cycles interrelated. NAND also supports faster read and write times and greater densities than NOR. NAND is therefore the primary solution used for storage of user data, but is not as commonly used for storage of code.

In fact, in 2016, roughly 66% of NAND flash (by units) was sold into data processing applications, with another 27% sold into wireless applications (predominantly handsets). The balance of the NAND market services embedded applications, primarily for data but also for code storage. Consumer, Automotive, Industrial, Communications and Other each had sub-5% share of the NAND market in 2016.

In terms of density, most NAND flash in 2016 was sold in densities of 64Gb and above, with 64Gb and 128Gb representing the majority of shipments. There is a constant push upward as data storage demand in consumer portables is far above what can be economically supplied with today's NAND chips and is continually rising as these devices become more complex.

The exhibits below display 2016 NAND flash revenue by application, alongside unit shipments by density.

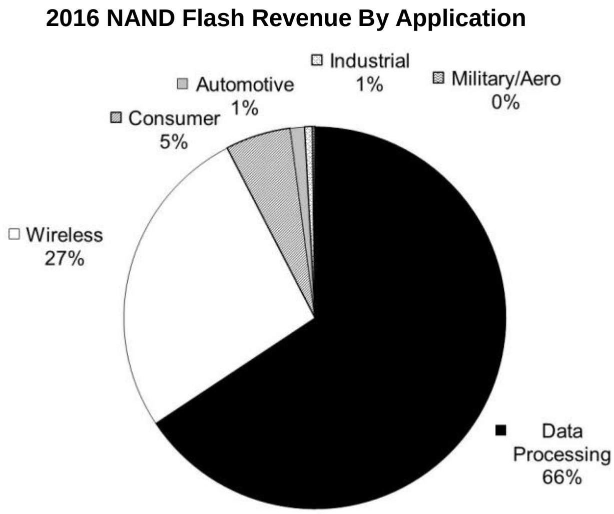

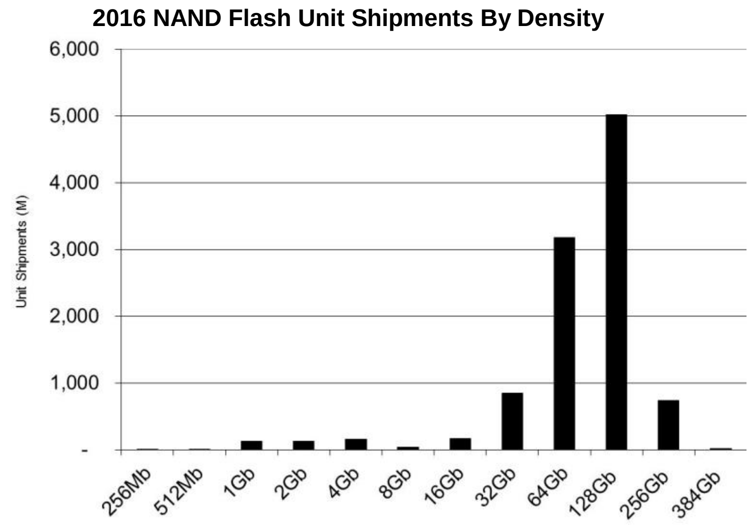

Source: Gartner, Oppenheimer & Co.

Market Forecast: The NAND market has grown steadily over the past few years, as a relatively small number of competitors served a dramatically growing need for solid state data storage, particularly in portable devices. Memory card demand is in the hundreds of millions of units, with wireless handsets and increasingly PCs driving most of the demand. A need for greater storage capacities and a profound preference for solid-state vs. HDD-based storage have also driven SSD demand. Only in the past couple of years has supply caught up with demand as new fabs have been built and new competitors have entered the market.

Looking ahead, NAND should benefit from continued strong demand in consumer electronics, primarily from embedded consumer applications like smartphones and tablets, but also increasingly from enterprise storage applications such as data centers and high-end SSD-enabled laptops. NAND is increasingly being utilized in PCs, both as a cache to boost performance and in solid state drives, which serve as a replacement for hard disk drives. Similar to the DRAM market, favorable pricing conditions and strong demand have led to significant NAND growth in 2017, with 1H17 NAND revenue up 49% Y/Y. This higher base driven by favorable pricing is the largest driver of our growth forecast. In total, we expect NAND to grow at a 9% CAGR, rising from $32.0 billion in 2016 to $48.9 billion in 2021.

Competitive Environment: The top NAND flash memory maker in 2016 was Samsung, with 38% revenue share. Toshiba was second with 20%, followed by Western Digital (through its SanDisk acquisition) with 15% and Micron with a 13% share. SK Hynix rounded out the top players with 11% share. Intel held 3% of the NAND market in 2016, but has been making a steady push into this market, primarily via its JV with Micron.

Legacy Memory: SRAM & NOR

SRAM (Static Random Access Memory) - A type of volatile memory that is configured in a flip-flop circuit, allowing cells to remain without being refreshed. A higher performance memory than DRAM, used mostly in communications applications and for high-speed data caching.

NOR Flash – Type of non-volatile rewritable memory supporting random access and execute-in-place. Used primarily for code storage.

SRAM Market Forecast

•We expect the SRAM market to decline at an 1% CAGR, from $532 million in 2016 to $518 million in 2021.

•Handsets have largely shifted from SRAM to DRAM as DRAM power consumption improves.

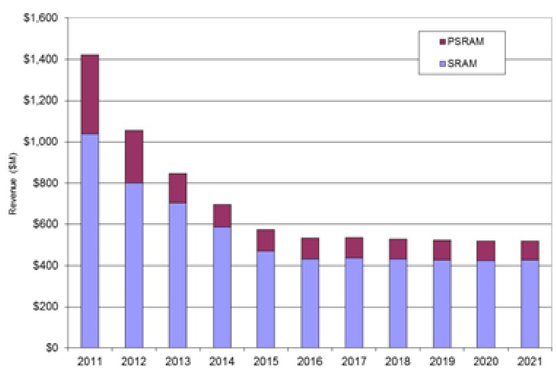

Source: WSTS, Oppenheimer & Co.

NOR Market Forecast

•We expect the NOR market to remain mostly flat at $1.7B from 2016 to 2021.

•As NAND flash has becoming increasingly competitive on a price/performance basis with NOR

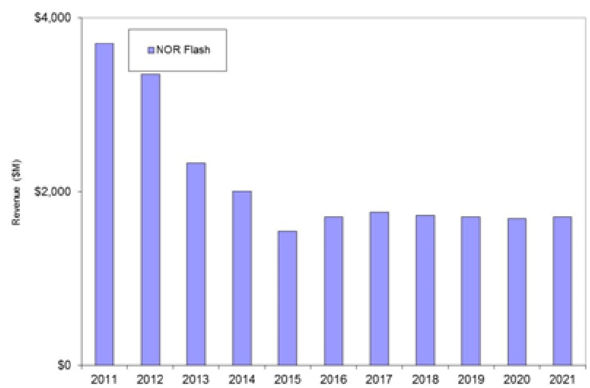

Source: WSTS, Oppenheimer & Co.

SRAM (static random access memory) is a type of volatile memory, which, like DRAM, loses stored information when powered down. SRAMs also support random access, which means that locations in memory can be written to or read in any order. Unlike DRAM, however, SRAM stores data bits in a flip-flop circuit using a minimum of four transistors, as opposed to the single transistor/single capacitor design used in DRAM. This allows current to flow through one side or the other based on which of the two transistors is activated, allowing for faster access time. The structure of SRAM means that it does not need to be refreshed, making it suitable for high-performance communications applications. It also consumes far less power than traditional DRAM, which has suited it well for the handset market.

Unlike DRAM, there is significant product sub-segmentation within SRAM beyond just density. Asynchronous SRAMs perform read and write operations sequentially, while synchronous SRAM overlaps reads and writes. Some SRAMs have multiple I/O ports or support advanced features such as ZBT (zero bus turnaround) or syncBurst (speeds up write operation), which make them ideal for high-end communications applications.

A segment of the SRAM market is actually not technically SRAM at all. PSRAM (or Pseudo SRAM) is actually DRAM modified to behave like SRAM. PSRAM has built-in refresh and address control circuits to mimic SRAM, but leverages the low-cost and high-density benefits of DRAM. These cost and density advantages have allowed PSRAM to displace standard SRAM in most handsets. Note that since PSRAM is sold into the same sockets as SRAM (primarily wireless), we include revenue and market share data here rather than in the discussion of DRAM.

SRAM finds heavy usage in the computing market as a cache for PC and server microprocessors. However, these SRAM cells are located directly on the CPU die; no discrete SRAM ICs are used.

NOR memory (5% of total flash revenue in 2016) is the most reliable and versatile flash memory, supporting random access and memory card program execution, known as “execute in place.” Writing and erasing can take several seconds, but reading is fast. For these reasons, NOR is primarily used for programming, code storage and other applications where reliability and security are a premium.

NAND flash (95% of total flash revenue in 2016) is cheaper than NOR and also somewhat less reliable. It does not support execute in place. It does, however, support faster read and write times and greater densities than NOR. NAND is therefore primarily used for data storage as opposed to code storage. NAND flash has increasingly been cannibalizing NOR flash in recent years as reliability has improved and storage density becomes more critical.

In terms of its architecture, NOR flash uses cells that resemble a standard transistor, except that they have two gates instead of just one. The presence of the second gate allows for electrons to become trapped within the cell, hence the ability to retain data when power is removed from the host device. First generations of NOR flash stored one information bit per memory cell in an architecture referred to as SLC, "single level cell." Advanced architectures permitting the storage of two or more bits per memory cell are referred to as MLC or "multi-level cell" architectures.

2016 NOR Flash Revenue By Application

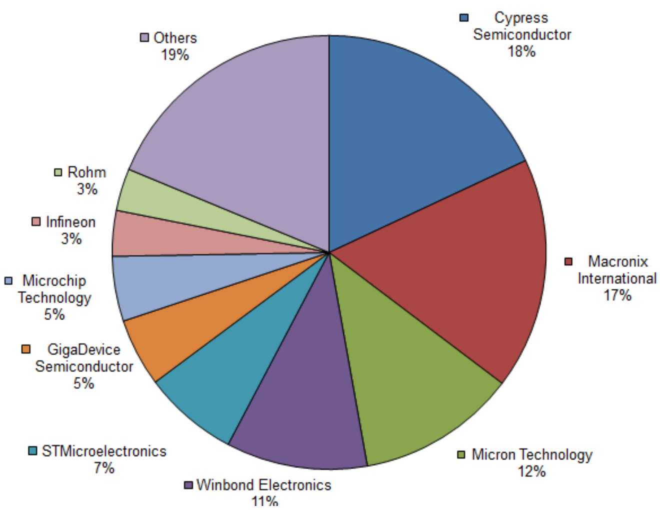

Source: Gartner, SIA, Oppenheimer & Co.

Competitive Environment: Cypress was the top vendor in 2016 of both NOR (through its acquisition of Spansion) and SRAM. Spansion was originally created as a flash manufacturing JV between AMD and Fujitsu known as FASL; in 2003 the two companies spun off the JV as a separate entity, contributing all aspects of operations including R&D, sales and marketing. Micron also manufactures both SRAM and NOR, holding ~12% share of the legacy memory segment. Macronix and Windbond supply primarily NOR in legacy memory, with 17% and 11% shares, respectively.

Discretes and Optoelectronics

Discretes:

Semiconductor devices that consist of a single transistor in a package

Three Key Markets:

Discretes

Sensors and Actuators

Optoelectronics

This section deals with all types of discretes, which are semiconductor devices consisting of a single transistor in a package (i.e., not an integrated circuit). These devices sell in extremely high volume (hundreds of billions per year) at extremely low ASPs (sometimes as low as a penny each). They are used in all types of electrical systems primarily as enabling devices—helping to move signals around the board or alternatively providing the bridge between the electronic system and the outside world. These devices are generally application agnostic, and are almost always sold through distribution. The three key types are:

Discretes: Discretes are semiconductor devices consisting of a single transistor in a package. They are often used to direct signals around a board and to control power flow in a system.

Sensors and Actuators: Sensors are specialized discrete semiconductor devices whose electrical properties can be translated into measurements of external stimuli. These include things like temperature, pressure, displacement, velocity, acceleration, stress, strain, or any other physical, chemical or biological property. Actuators are the opposite of sensors; actuators create mechanical motion by converting various electrical signals to rotating or linear mechanical energy.

Optoelectronics: Optoelectronics are semiconductor devices designed to produce and receive light waves. This includes displays, lamps, couplers and other opto-sensing and emitting semiconductor devices. Note that sensors that measure light and capture images are grouped with optoelectronics.

Discretes

Discretes - Discretes are semiconductor devices consisting of a single transistor in a package. They are used as enabling devices within larger systems, and are generally application agnostic.

Key Types: Diodes, small signal transistors, switching transistors, power transistors, rectifiers (power diodes), thyristors.

Market Forecast

- We expect the market for discretes to grow at a 3% CAGR, from $18.2 billion in 2016 to $21.1 billion in 2021.
- Growth for these multi-market devices should slightly trail the broader industry as ASSPs seek to integrate discretes in some high volume applications.

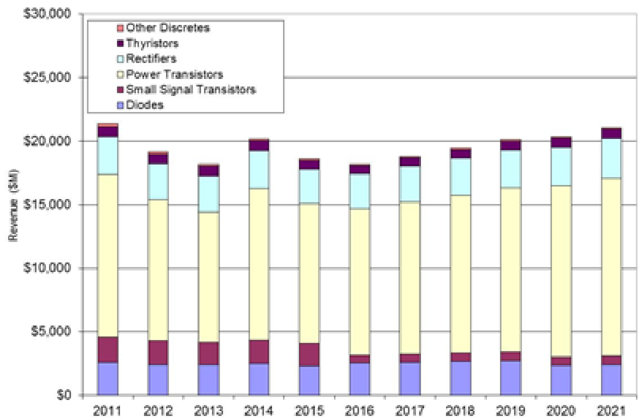

Source: WSTS, Oppenheimer & Co.

Key Competitors

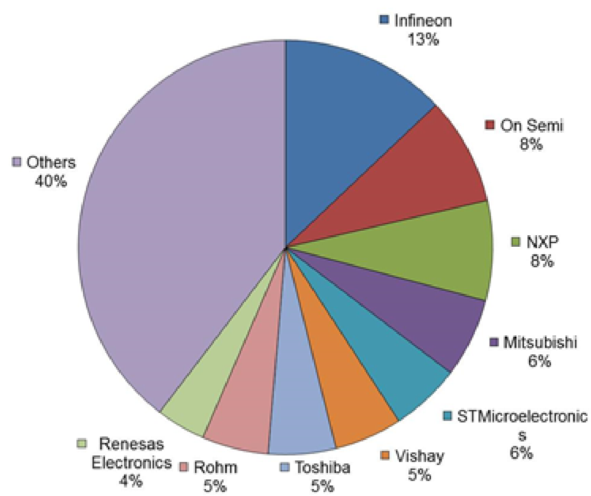

Source: Gartner, Oppenheimer & Co.

Discretes are semiconductor devices consisting of a single transistor in a package. These devices sell in extremely high volume (hundreds of billions per year) at extremely low ASPs. They are used in all types of electrical systems primarily as enabling devices—helping to move signals around the board as well as to control power flow—and are generally application agnostic. There are several major categories of discrete devices, including:

Diodes: Devices that act as a one-way valve. Key types include small signal diodes, zener diodes, transient protection diodes and RF/microwave diodes.

Small Signal and Switching Transistors: Transistors with power dissipation of less than 1W.

Power Transistors: Transistors with power dissipation of more than 1W. Key types include RF/microwave power transistors, bipolar general purpose transistors and insulated gate bipolar transistors (IGBTs) and modules.

Rectifiers (Power Diodes): Devices used to convert AC (alternating current) to DC (direct current). These devices are classified by the amount of current supported.

Thyristors: Devices that can turn on when activated by a gate signal.

Other Discretes: These include varactor tuning diodes, selenium rectifiers and other polycrystalline devices.

On a revenue basis, more than half of discrete sales are power transistors, at 63% of the total in 2016. Rectifiers, diodes and small signal transistors were 15%, 14%, and 4%, respectively. Thyristors and other discretes represent the remaining 4% of the market.

In terms of application, discretes go into every electronic device produced. In 2016, the revenue split was 26% automotive, 27% industrial, 17% communications, 14% consumer, 12% data processing and 4% military/aerospace.

The exhibits below display discrete semiconductor revenue share by product type and application.

2016 Discrete Revenue By Product Type

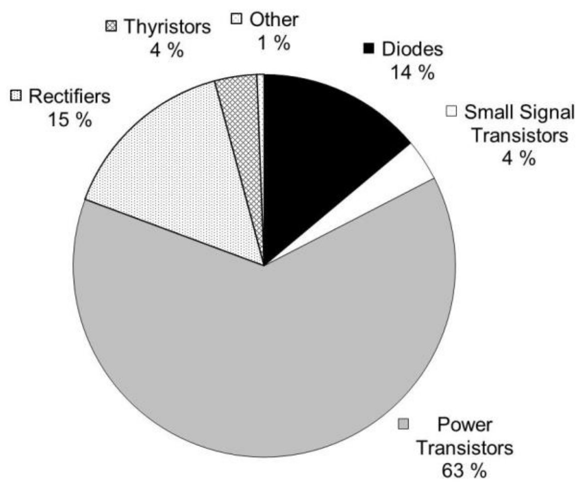

2016 Discrete Revenue By Application

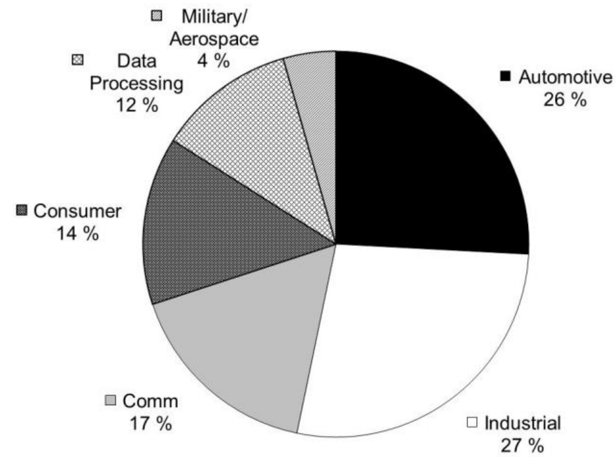

Source: WSTS, Gartner, Oppenheimer & Co.

***Market Forecast*:** Discretes sell into every electronic application and are therefore tied to the semiconductor cycle and to demand for electronics in general. Many of these products are commodities and thus are subject to rapid changes in pricing as fab capacity rises and falls with the semiconductor cycle. Countering this effect, however, is the fact that discretes do not see the same concentration in compute-intensive digital equipment like PCs and handsets, which causes a lot of volatility for the big dollar segments such as microcomponents, digital logic, and memories. Discretes therefore tend to have more stable unit demand than the semiconductor market as a whole. At the same time, they tend to "miss out" on the growth generated by those advanced products.

Overall, we expect the market to grow from $18.2 billion in 2016 to $21.1 billion in 2021, a CAGR of 3%.

***Competitive Environment*:** The market for discretes is large and fragmented. The largest vendors in 2016 were Infineon (13% share), On Semi (8%), NXP (8%), Mitsubishi (6%), and STMicro (6%). Other large vendors include: Vishay, Toshiba, Rohm, and Renesas.

Sensors and Actuators

Sensors – Specialized discrete semiconductor devices that translate real world input into electrical signals.
Actuators - Actuators are the opposite of sensors, taking electrical signals and translating them into some form of real world output.
Key Types: Temperature sensors, pressure sensors, acceleration and yaw rate sensors, magnetic field sensors, LEDs, motors, solenoids, speakers.

Market Forecast

•We expect the market for sensors and actuators to grow at a 6% CAGR, rising from $10.1 billion in 2016 to $13.8 billion in 2021.

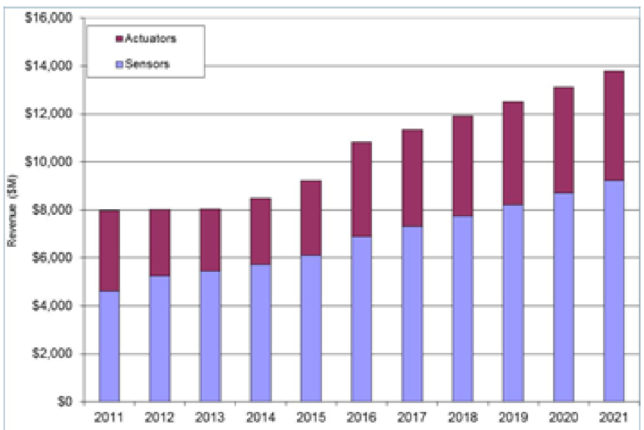

Source: WSTS, Oppenheimer & Co.

Key Competitors

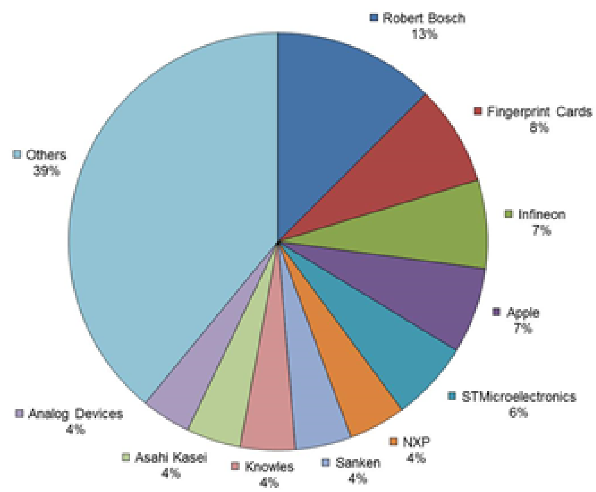

Note: Market share data is for sensors, and excludes actuators.
Source: Gartner, Oppenheimer & Co.

Sensors are specialized discrete semiconductor devices whose electrical properties can be translated into measurements of external stimuli. These include things like temperature, pressure, displacement, velocity, stress, strain, acceleration, or virtually any other physical, chemical or biological property. The category also includes actuators, which translate signals into movement, acting as the opposite of a sensor (the mechanisms that move the head assembly on a disk drive or an arm of a robot are two examples of actuators).

In 2016, actuators represented 39% of total sensor and actuator revenue, while sensors rounded out the remaining 61%. Within sensors, magnetic field sensors represented 23% of the total; inertial were also 23%; fingerprint sensors were 22%; pressure sensors were 14%; MEMS microphones were 10%; and other sensors accounted for 7%. Note that image sensors are not included in this category and are instead grouped with optoelectronics.

In 2016 by end market, Comms were 42% of the sensor market; automotive 39%; data processing 7%; industrial 5%, consumer 5%, and Military/Aerospace were 1%.

***Market Forecast*:** The market for sensors and actuators is expected to grow at a 6% CAGR over the next five years, from $10.1 billion in 2016 to $13.8 billion in 2021. The growing digitization of devices, especially in automobiles, as well as the move to silicon-based sensors from other materials is expected to drive this growth.

***Competitive Environment*:** The market for sensors is smaller and thus less fragmented than for general discretes, but still many niche players have been able to build defensible positions. The largest vendor in 2016 was Robert Bosch with 13%, followed by Fingerprint Cards with 8%. Infineon held 7%, Apple held 7% and STMicro also held 6%. Other suppliers include NXP, Sanken, Knowles, Asahi, and Analog Devices. Note that this market share data does not include a breakout for actuators (data for this segment is not available).

2016 Non-Optical Sensor Revenue By Product Type

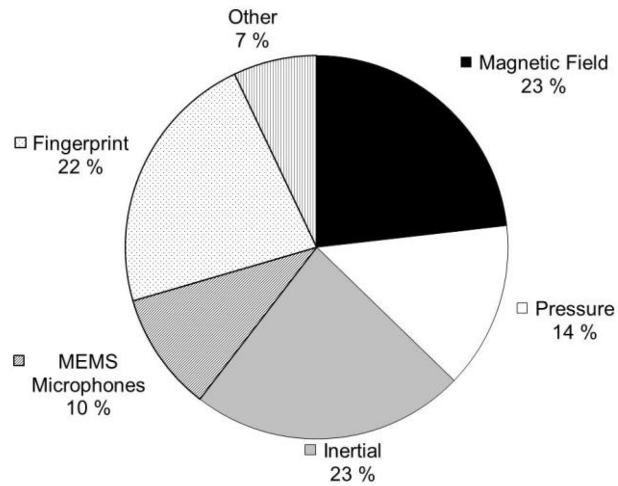

2016 Non-Optical Sensor Revenue By Application

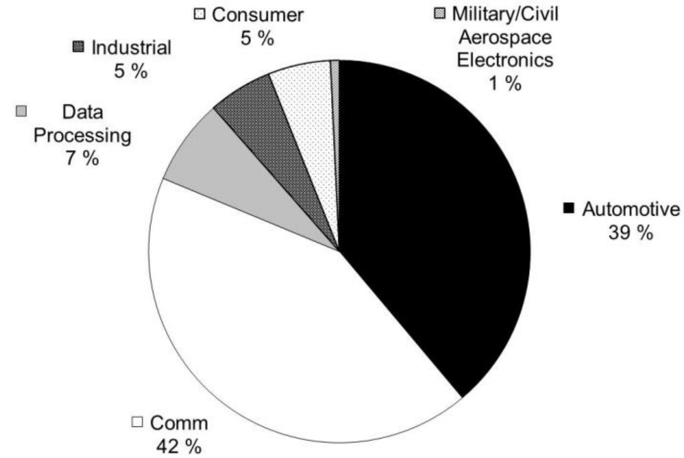

Source: Gartner, Oppenheimer & Co.

Optoelectronics

Optoelectronics - Semiconductor devices designed to produce and receive light waves, including displays, lamps, couplers, image sensors, and other opto-sensing and emitting devices.
Key Types: Displays, lamps, couplers, laser devices, image sensors (CCD, CPD, SSP, CMOS), infrared devices.

Market Forecast

•We expect optoelectronics to grow at an 5% CAGR, rising from $33.2 billion in 2016 to $42.4 billion in 2021.

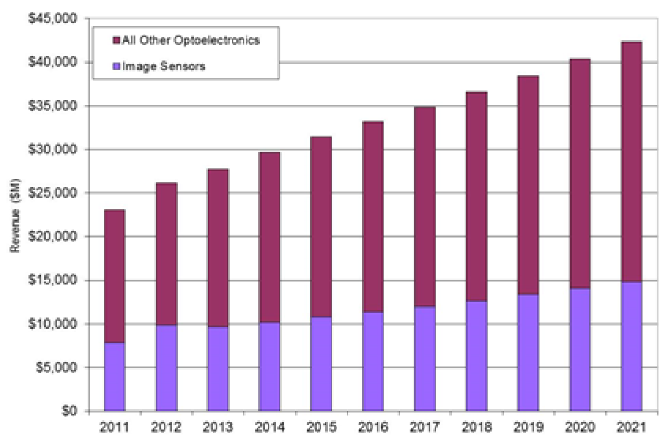

Source: WSTS, Oppenheimer & Co.

Key Competitors

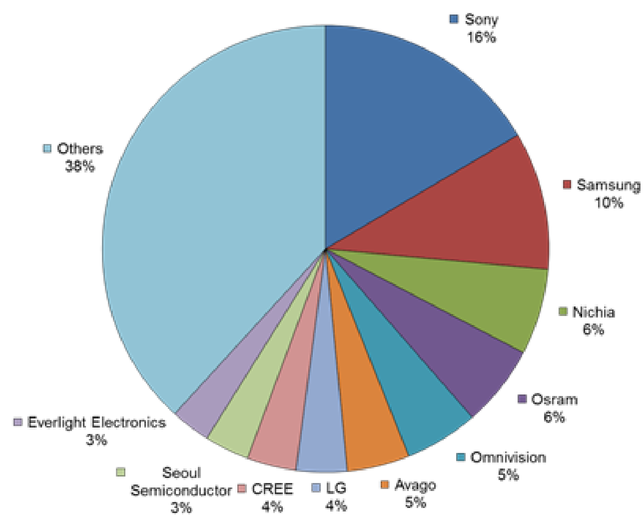

Source: Gartner, Oppenheimer & Co.

Optoelectronics are semiconductor devices designed to produce and receive light waves. This includes displays, lamps, couplers and other opto-sensing and emitting semiconductor devices. Only semiconductor devices are included here, however; the category excludes liquid crystal devices and displays, incandescent lamps and displays and other non-semiconductor optoelectronics components.

The best-known type of optoelectronic device is the LED (light emitting diode), a diode that emits light when charged. LEDs are housed in a small "bulb" and are used in all kinds of readouts (see diagram). Even traffic lights have begun to use LEDs instead of less efficient, traditional filament bulbs.

About 34% of optoelectronics revenue is actually derived from image sensors, including CCDs and CMOS image sensors. These optoelectronics sensors perform image pickup for digital cameras and camera phones.

.

Light Emitting Diode

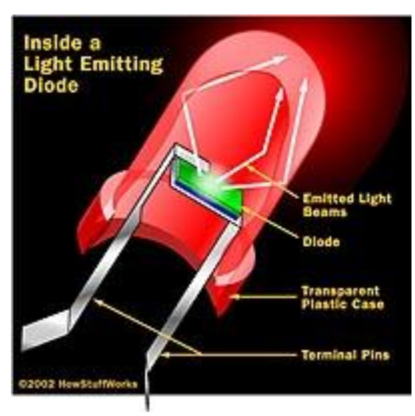

Source: I. Weiss, Oppenheimer & Co.

Like discretes, optoelectronics are application agnostic, though they are more heavily tied to the communications and industrial markets, which were 31% and 33% of sales in 2016, respectively. 14% of sales were derived from the consumer market, 11% from data processing and the remaining 12% from automotive and military/aerospace.

The exhibits below display optoelectronics revenue by device type and application.

2016 Optoelectronics Revenue By Product Type

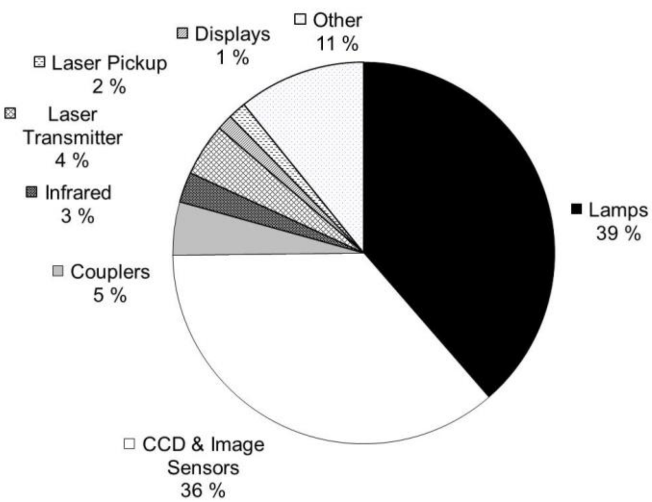

2016 Optoelectronics Revenue By Application

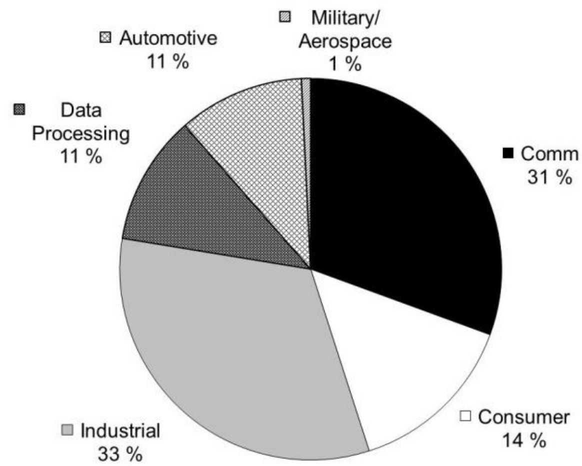

Source: WSTS, Gartner, Oppenheimer & Co.

***Market Forecast*:** The market for optoelectronics is expected to slightly outpace the non-memory semiconductor industry through 2021, growing at a 5% rate. After several years of rapid growth driven by digital cameras and smartphones, we expect image sensors to see a more modest CAGR as pricing pressure partially offsets continued unit growth. Overall, optoelectronics revenue is expected to grow from $33.2 billion in 2014 to $42.4 billion in 2019. We expect autonomous driving and artificial intelligence to serve as a demand driver for image sensors as data collection becomes increasingly important to their success.

***Competitive Environment*:** The market for optoelectronics is large and highly fragmented. The top producer in 2016 was Sony with a 18% market share, derived from its No. 1 position in image sensors (47% of market). Samsung was the No. 2 vendor with a 10% share, and held the No. 2 supplier of image sensors (18%) and other (6%).

The exhibits below display market share data for image sensors and all other optoelectronics.

2016 Image Sensor Market Share

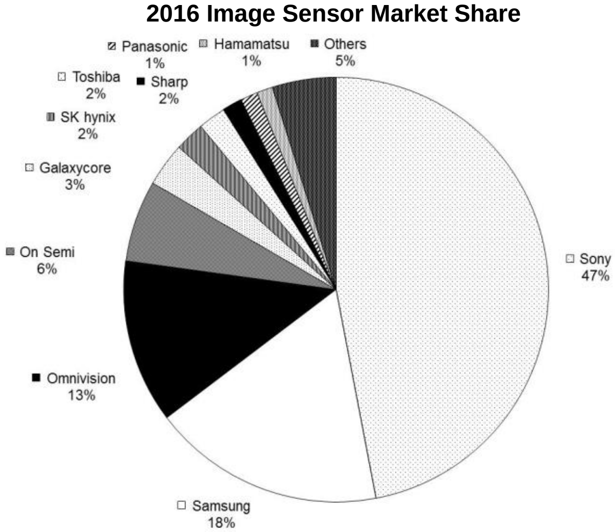

2016 All Other Optoelectronics Market Share

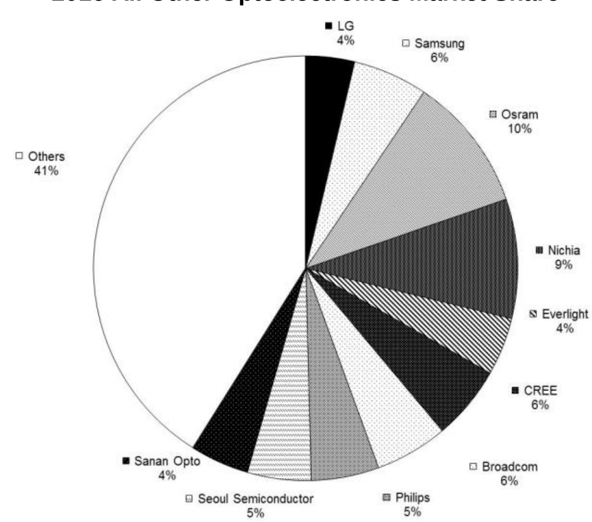

Source: Gartner, Oppenheimer & Co.

Section V: End Market Summary

This section discusses the major semiconductor end markets

For each end market we discuss:

Key technologies & end market fundamentals

End market unit forecast

Semiconductor revenue forecast

Semiconductor block diagram

Competitive environment

Semiconductor trends

End Market Summary

The semiconductor sector serves a number of important end markets

Six Key End Markets:

Computing

Wireless

Networking

Consumer

Telecom/ Datacom

Automotive

We now move from a discussion of semiconductor devices to the electronic systems that use them. We discuss several key equipment types grouped into six end market categories, and for each type of equipment, we discuss the end market unit forecast, our semiconductor revenue forecast, the competitive environment for both equipment and ICs, and market and technology trends. In general, we limit the equipment types we discuss to those that are implemented primarily using application specific standard and/or custom ICs. There are a multitude of electronic devices that are implemented using mostly generic components—e.g., analog SLICs, microcontrollers, general purpose logic, discretes, etc.—which, while they may ship in significant volume, do not present themselves as opportunities for specialization. We break down the key semiconductor end markets into the following categories:

- **Computing:** Equipment used in personal computing and data processing, including PCs and servers, PC displays, hard disk drives, optical storage, and printers/multi-function peripherals.
- **Networking:** Client and infrastructure equipment for local area, wireless, personal area and storage networks.
- **Telecom/Datacom:** Service provider and enterprise equipment used to implement voice and data networks.
- **Wireless:** Cellular handsets and wireless infrastructure equipment.
- **Consumer:** Digital consumer devices, such as set-top boxes, DVD players, MP3 players, game consoles, etc.
- **Automotive:** Electronic systems and subsystems used in automobiles.

Block Diagram Example: A Basic System

In the sections that follow, we will use block diagrams to identify key semiconductor opportunities.
A typical electrical system block diagram looks like this:

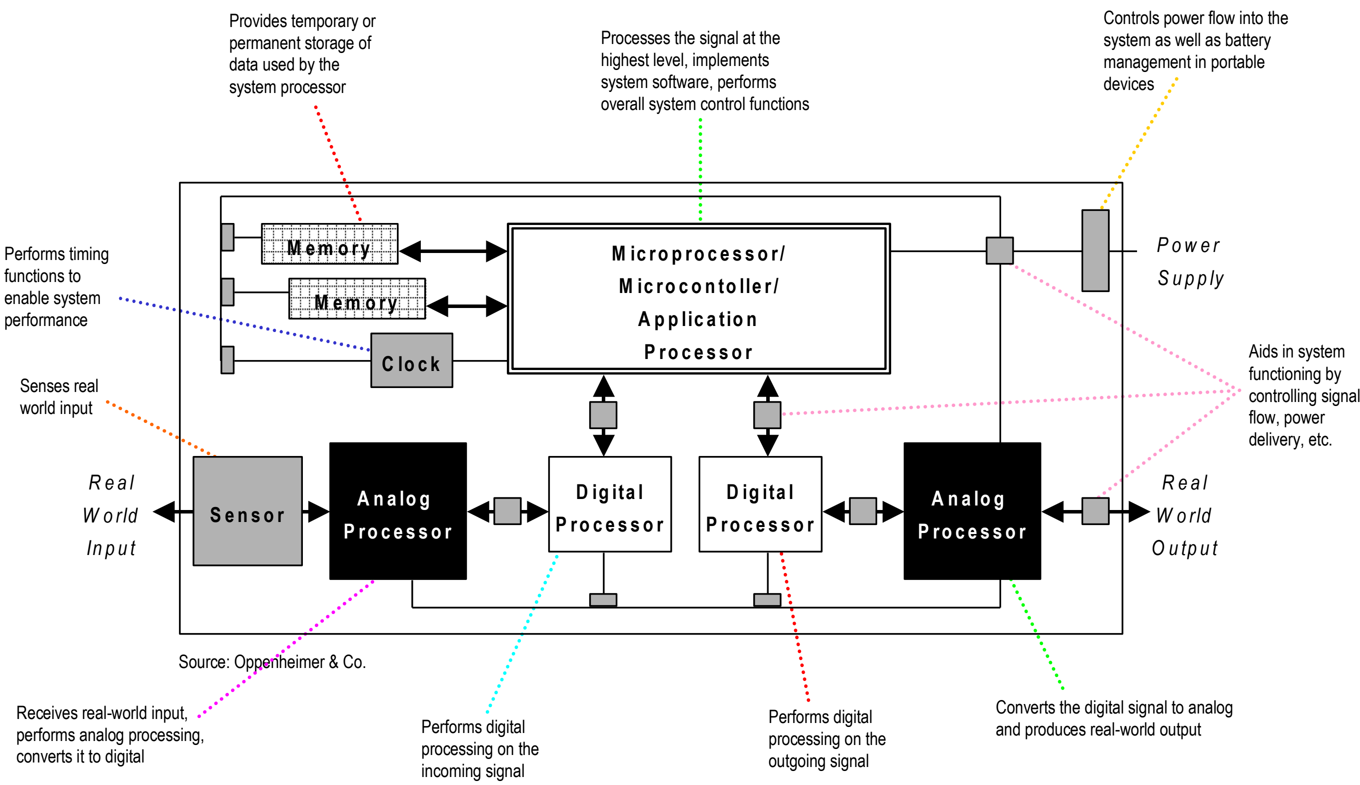

Source: Oppenheimer & Co.

The diagram above, which we showed earlier in Section III when we introduced the different types of semiconductor devices, displays a generic electronic system. In the sections that follow, we show application-specific electronic systems using the same format. We will use black to denote analog devices and white to denote digital devices. Double borders denote microcomponents, while a checkered pattern denotes memory. Note that in most cases, we will not show the multitude of enabling devices (gray) that populate the board.

Computing

Three Key Markets:

PCs, Tablets and Servers

PC Displays

Hard Disk Drives / Solid State Drives

Computing remains the largest single application for semiconductors and includes equipment such as desktop PCs, notebooks, tablets, servers, displays, hard disk drives, printers and other peripherals. Although these devices are pervasive in the United States, Western Europe and Japan, emerging markets are major growth opportunities. In discussing the end markets for computing semiconductors, we focus on three key product areas:

- **PCs and Servers:** Includes all semiconductor devices—including memory—that are used on desktop and notebook motherboards and in industry-standard servers. This area excludes semiconductors sold into subsystems and peripherals, such as hard disk and other mass storage devices, displays, printers and other peripherals. We also provide a tablet unit forecast. Each of these devices are discussed separately.
- **PC Displays:** Includes all displays designed specifically for use with PCs, including CRT, desktop LCD, and notebook LCD screens.
- **Hard Disk Drives:** All hard disk drives used in PCs and servers, enterprise storage, and consumer applications.

PCs and Servers

Despite an ongoing shift to mobile computing, PCs remain a large semiconductor market and a huge portion of technology spending. We believe the age of the PC is behind us, and model PC units in secular decline. Server growth remains healthy.

PC Unit Shipments and YoY Growth

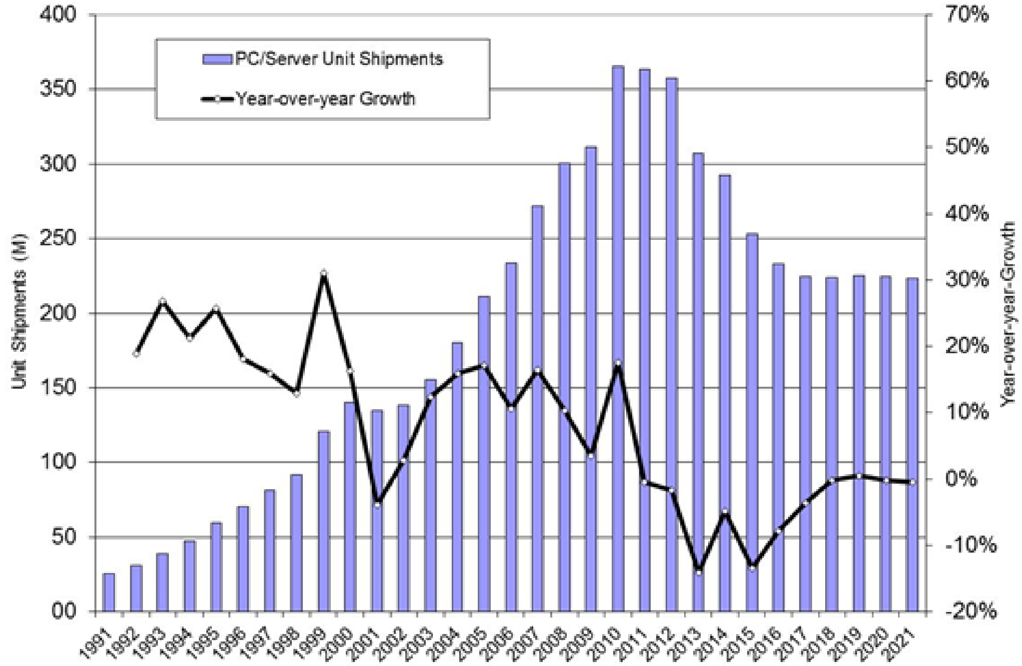

Source: IDC, Oppenheimer & Co.

PCs remain at the center of the technology universe and are a huge portion of IT and home technology spending. The market has matured a great deal over the past decade, with the commoditization of PCs forcing consolidation in the industry and average selling prices declining rapidly as system differentiation diminishes. The PC market was beginning to reach saturation even before the onset of smartphones and tablets. The rise of these cheaper and more portable computing form factors has dramatically altered the PC industry. We believe that the total PC market is now in a state of secular decline. Units declined another 13% in 2015 and 8% in 2016, and have stabilized somewhat thus far in 2017.

PCs and Servers

Servers have seen the most consistent growth, while both desktop and notebook units appear to have reached market maturity/saturation. Computing has gone mobile with the rise of smartphones and longer replacement cycles. Notebook unit declines have declined since 2012

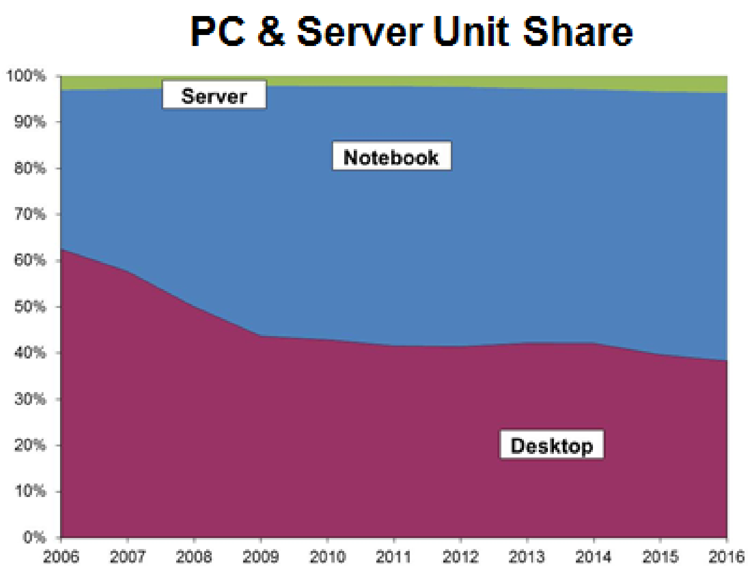

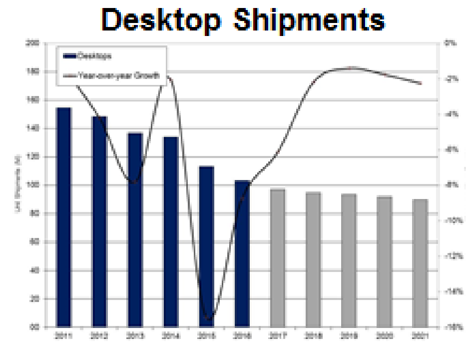

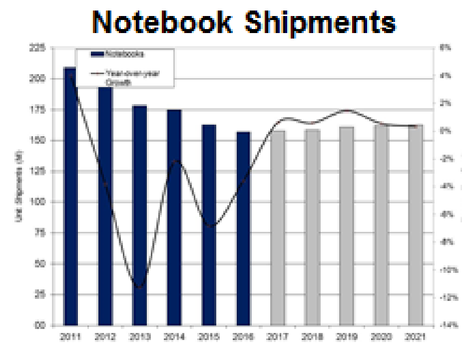

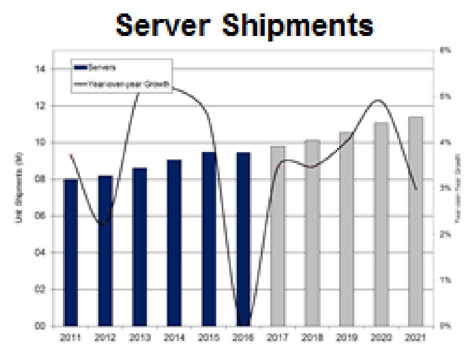

Source (all graphs): IDC, Oppenheimer & Co.

By platform, desktops have historically been the largest market. In 2009, notebooks surpassed desktops as the largest segment of the PC market. Today, desktops represent 40% of total units shipped. Desktops are also the most mature market, having declined at a 10% CAGR over the past five years vs. slightly larger than the overall market. In general, growth in desktops is limited to emerging markets, high-end consumer gaming and developed markets enterprise refreshes.

Notebook growth had been accelerating until 2012; however, as penetration rates in developed markets peaked, sales have declined, and units have also declined at 9% CAGR over the past five years, with sharp declines in 2013 and 2015. Notebooks now represent 56% of shipments. Notebooks have been the most affected part of the market by the emergence of smartphones and tablets. We believe notebook units are likely to decline slightly as consumers juggle notebook/smartphone/tablet purchase decisions.

x86 servers have been the comparative outperformer, having grown 4% over the past five years—though still small at 4% of total unit shipments. Server OEMs have done an excellent job of packing more performance and functionality and improving ease-of-use into low-end and mid-range servers over the past few years, opening the market to more small and medium-sized businesses, which have historically depended on either custom-built systems or simple small-business software running on a normal desktop PC. In recent years the prevalence of virtualization and hyperscale datacenters/cloud computing has also driven growth. x86 servers are also cannibalizing the mid- and high-end proprietary server market, offering cost and configurability advantages over proprietary systems from OEMs like IBM and HP. The growth of the microserver market based on alternative SoC architecture has yet to disrupt the traditional x86 ecosystem.

PCs and Servers

PC shipments are shifting from the mature markets of North America, Western Europe, and Japan to developing economies in Asia, Eastern Europe, and Latin America. The distribution channel is growing in importance, serving local consumers in these fast-growing emerging markets.

PC Unit Shipment Share By Region

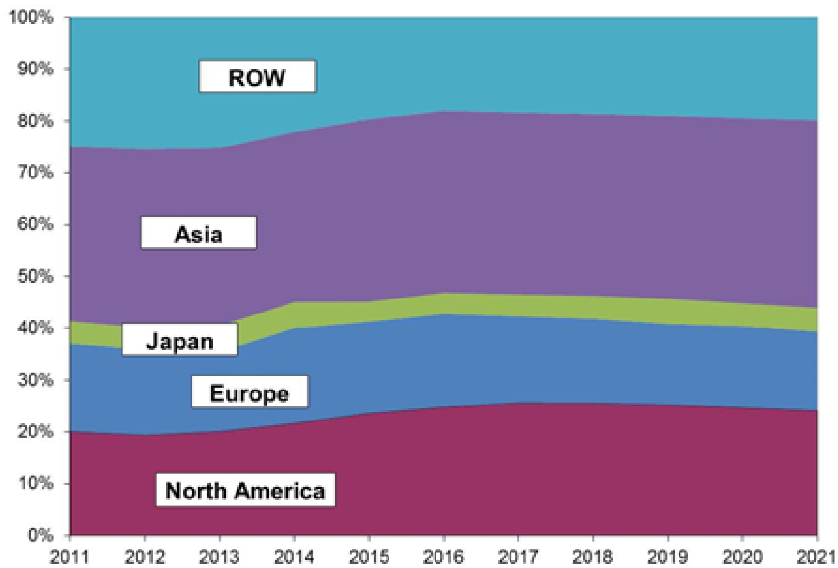

Source: IDC, Oppenheimer & Co.

Looking at the market by geography, PCs have shown a sustained shift away from mature markets and toward developing economies, especially in Asia. North America was overtaken as the largest market for PCs in 2007 by Asia, and North America now represents just 25% of total shipments (as of 2016). This market has declined at a 3% CAGR the past five years. Western Europe—18% of PC shipments—has declined at a 6% CAGR over the same period.

Japan, Asia and other emerging markets, have also declined in recent years. Unit shipments have declined at a 6% CAGR over the past five years in Asia, an 8% decline in Japan, and a 13% CAGR in the “Rest of World” region. As a percentage of total shipments, Asia and rest of world account for 39% and 18% of unit shipments, respectively.

PCs and Servers

Desktop and Notebook PC Unit Forecast

- Total PCs declining 1%, from 223 million units in 2016 to 212 million units in 2021.
- Desktops declining 3%, from 93 million units in 2016 to 80 million units in 2021.
- Notebooks roughly flat at ~130 million units
- Notebooks should decline slightly in emerging markets (ex US which should remain flat) and remain relatively flat in emerging markets.

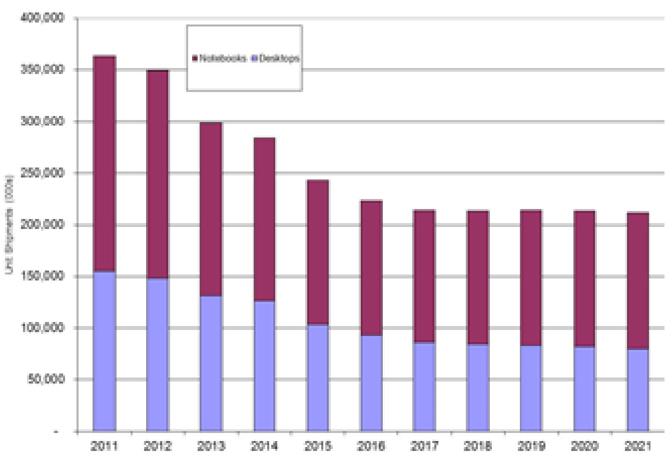

Source: IDC, Gartner, Oppenheimer & Co.

Desktop and Notebook PC Suppliers

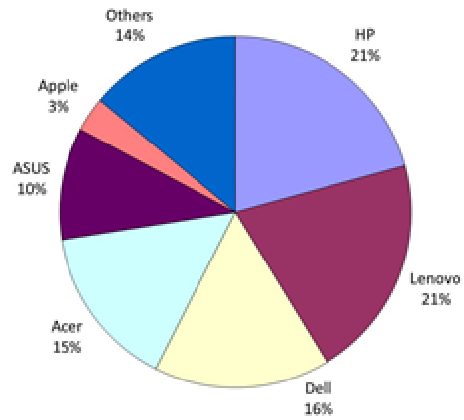

Desktops	***Notebooks***
Lenovo	Lenovo
HP	HP
Dell	Dell
Acer	Acer
Apple	ASUS
ASUS	Apple
	Toshiba
	Samsung

Source: IDC, Oppenheimer & Co.

PC Unit Forecast: The past several years have seen a rapid transformation in the PC market. Notebook computing has been relegated largely to "content creation" as many consumers have shifted "content consumption" to more mobile smartphones and tablets. The notebook market suffered its first ever annual decline in 2012, and continued declining since. The desktop market has matured, with growth limited to gaming, and enterprise refresh.

Going forward, we expect PC unit growth to decline at a 3% CAGR, declining from 223 million units in 2016 to 212 million units in 2021. We expect the worst of the declines are behind us, and expect more moderate declines through our forecast period. The risk to our forecast is likely to the downside, not to the upside. We expect the trend away from desktops to continue, declining through 2021; overall we expect desktop PCs to decrease from 93 million units in 2016 to 80 million in 2021. Note that growth opportunities remain in high-end desktop gaming and pockets of emerging markets.

We expect the rapid consumer shift to cheaper and more mobile smartphones/tablets to continue to erode the notebook market. We expect total notebook units to remain roughly flat at 130 million units from 2016-2021. While units could remain flat, but it is likely that ASPs will fall under pressure. Notebooks are no longer the computing platform of choice and consistently downward pressure on smartphones/tablets is likely to further pressure notebook pricing. We do not see Ultrabooks, convertible notebook/tablets or hybrid devices as a possible pocket of growth.

We believe most PC vendors have come to terms with the lack of elasticity in the market; i.e., their price cuts can be successful in winning market share but have been less than effective in growing unit demand for the industry as a whole. We also believe the market is sufficiently consolidated. We would anticipate most of the major PC OEMs to de-emphasize the PC business over time with an increased focus on smartphones and tablets.

The exhibits below display our unit shipment forecast for desktops and notebooks.

Desktop PC Unit Shipments, 2011-2021E

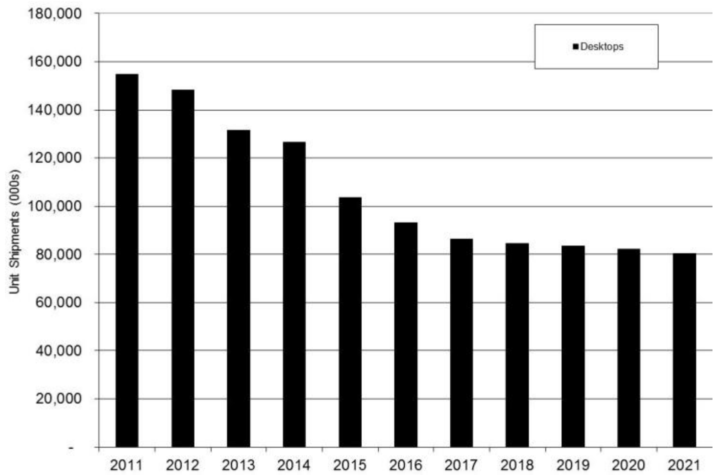

Notebook PC Unit Shipments, 2011-2021E

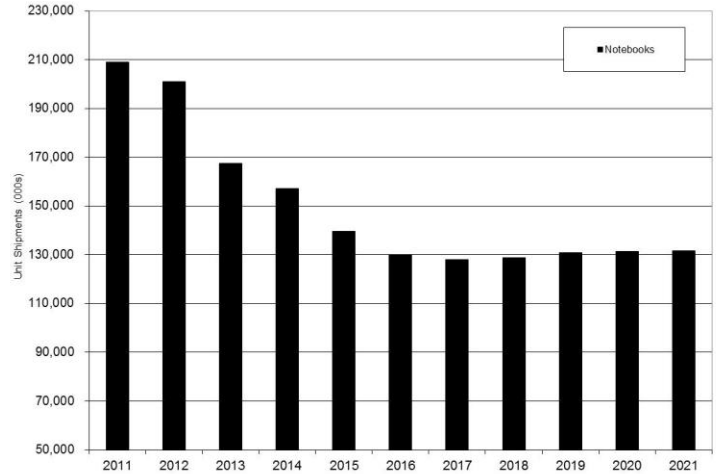

Source: IDC, Gartner, Oppenheimer & Co.

PC Suppliers: The PC market is served by several large double-digit share-holders: HP, Lenovo, Dell, Acer and ASUS as well as a few notebook specialists. In total, the top six vendors comprised 86% of the PC market in 2016. In 2016, HP and Lenovo held 21% unit share with Dell's 16%, Acer's 15%, and ASUS' 10% rounding out the top five.

PCs and Servers

Tablet Unit Forecast

•We model the total tablet market to decline at a 3% annual rate, from 175 million units in 2016 to 153 million units in 2021.

•Android become the largest tablet operating system in 2013 and we expect this to remain true through 2021. Much of Android's growth has come from low-cost white box tablets.

•iOS remains a dominant platform in tablet computing. We expect Android and iOS will account for 80%-plus of the market in 2019.

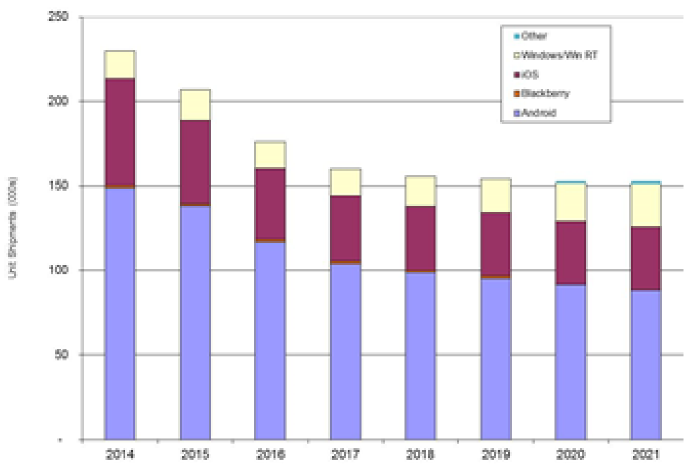

Source All: IDC, Oppenheimer & Co.

Tablet Suppliers

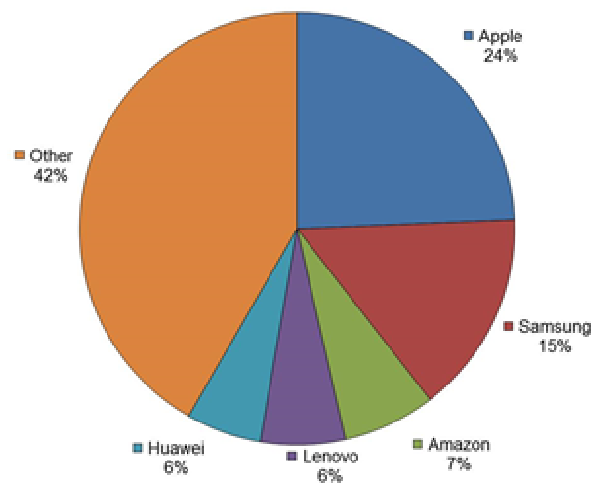

Tablet Unit Forecast: Apple has maintained a dominant share of the tablet market since the launch of the iPad, however Android has emerged as a viable competitor. While Apple likely maintains its number 1 overall market position, the explosion of low end white-box tablets in emerging markets and Androids lower cost alternatives present meaningful competition for units. At the high end of the market, differentiation and seamless OS interactivity with other devices (handset, TV etc.) will remain of primary importance. Apple aside, Samsung and Amazon have emerged as the two most formidable OEM competitors, however the market remains fragmented.

We expect the tablet market to decline from 153 million units in 2016 to 153 million units in 2021, a 3% CAGR. We believe that much of the growth of tablets is being stymied by the increasing screen size and functionality of smartphones, limiting the use case of tablets.

PCs and Servers

Server Unit Forecast

- Total servers growing at a 3% CAGR, from 9.7 million units in 2016 to 11.5 million units in 2019.
- x86 servers represent 99% of volume today.
- The most notable trend in the server industry is a change in the demands being placed on the data center. There is an increasing focus on cost/efficiency.
- While ARM remains a threat, the increased buying power of hyperscale cloud datacenters has yet to spur widespread traction

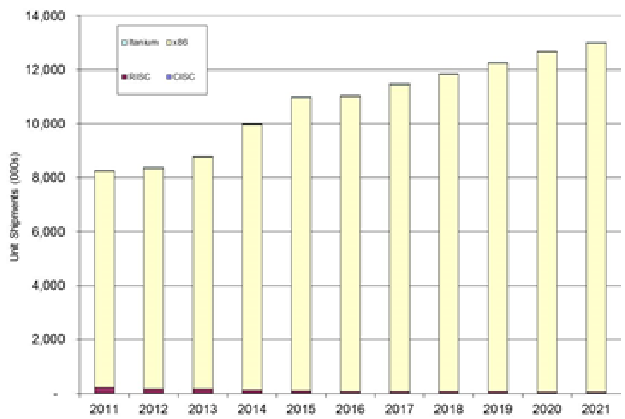

Source: IDC, Oppenheimer & Co.

Server Suppliers

Industry Standard Servers (x86)

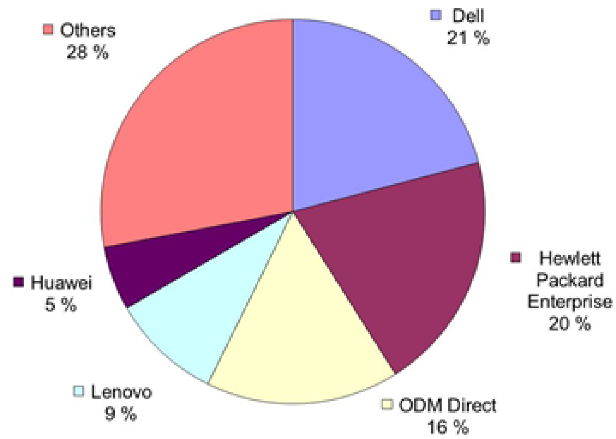

Source: IDC, Oppenheimer & Co.

Server Unit Forecast: As of 2016 industry standard servers are 99% of unit shipments. For some time, ARM has represented a low-power, efficient alternative to x86; however, it has yet to take material share from traditional vendors with established ecosystems and a track record of workload success.

Looking ahead, we expect a 3% CAGR for server units, with the market growing from 9.7 million units in 2016 to 11.5 million units in 2021. x86 servers should continue to dominate the market in terms of units.

Server Suppliers: The x86 server market is served by several key players, though each of their market shares has been dramatically reduced over the past few years: HP leads the market with 20% unit share in 2016 (down from 37% in 2010), while Lenovo holds 9% share. The rest of the market remains fragmented, with dramatic growth in the white box server market as customers increasingly cut out OEM services to build directly custom servers at the ODM level. We expect ODM Direct sales to continue to account for a larger share of the market as hyperscale datacenter customers gain share from traditional enterprise.

PCs and Servers

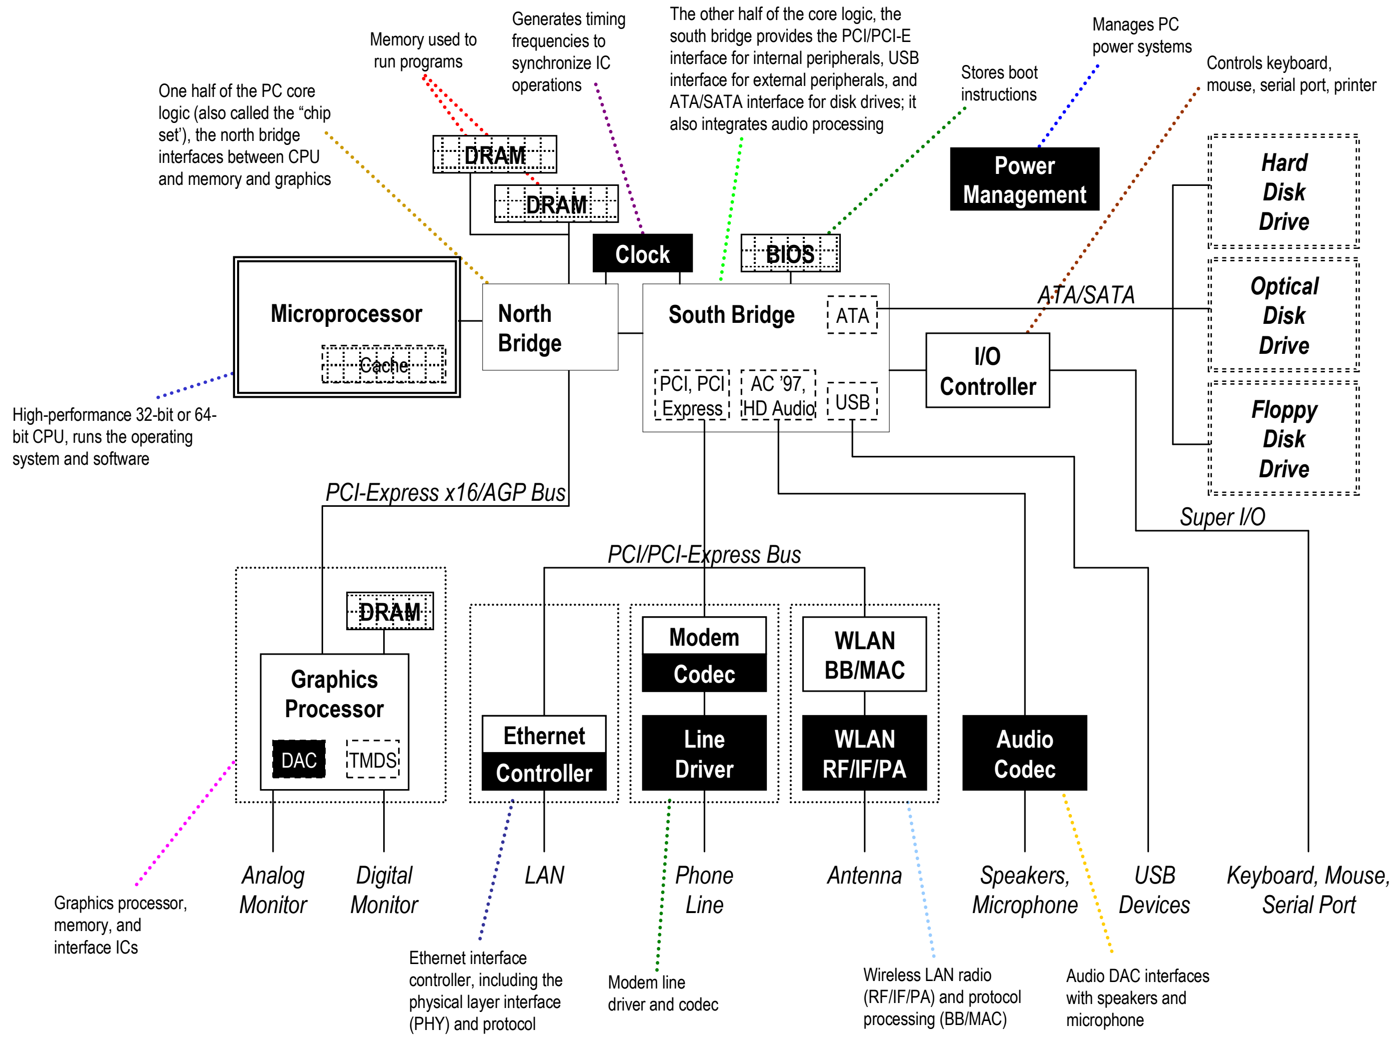

Source: Oppenheimer & Co.

Note: The diagram above shows a typical full-featured desktop PC. Notebook PCs contain some additional semiconductor content, including battery management and PC card controllers. Servers look slightly different in that they can have multiple processors, more memory, larger and less integrated chipsets and less multimedia content.

The nature of computing requires a variety of powerful devices with diverse functionality to handle a wide set of data types. The easiest way to think of a PC is to draw an analogy to the human body. The CPU (central processing unit), a microprocessor, performs all high-level processing, essentially the high-level brain functions of critical thinking and problem solving. RAM (random access memory) acts as short-term memory, used by the CPU as a workspace to process its instructions. BIOS, another form of memory, stores boot instructions (similar to autonomic nerves), while a clock provides synchronization functions, keeping all systems in check. The chipset (includes north bridge and south bridge) acts as the base of the spinal cord, relaying instructions to and from the CPU and connecting it to the rest of the PC over the system bus (usually PCI or PCI-Express), which functions as the spinal column. The hard disk, or solid state drive provide long-term storage, which is retrieved by the CPU and moved into RAM when needed, similar to a human being's long-term memory. Together, these systems function as a cohesive central nervous system, powerful enough to perform complex computations and versatile enough to handle a wide variety of inputs.

Outside these core systems, various peripheral devices perform the specific input and output processing to connect the PC to the outside world. The I/O controller interfaces with simple input devices such as the keyboard, mouse, parallel and serial ports (used by older, pre-USB printers and other peripherals). Graphics processor and interface chips produce the visual image displayed on the monitor (allowing the computer to draw pictures for display), while audio controllers and codecs produce sound (allowing the computer to literally speak and be spoken to). Ethernet, modem and wireless LAN/Bluetooth ICs connect the PC with other data sources (analogous to a person placing a phone call) and USB/Thunderbolt ports interface with additional peripherals (e.g., scanners, printers, smartphones).

PC and Server Semiconductor Device Summary

Processing

- **Microprocessor:** The central processing unit (CPU) in a PC is a high-performance 32- or 64-bit microprocessor. It performs all major processing of software code and communicates with all other devices in the system. It uses a modified CISC (Complex Instruction Set Code) architecture to process the lengthy sets of instructions used in typical computer software. Aside from this compute "core," the CPU usually includes on-chip memory called a "cache" as well as additional logic "extensions," which are embedded instructions designed to optimize the execution of certain types of critical functions (i.e., multimedia, security). Most mainstream CPUs now use multiple cores, usually two or sometimes even four. Microprocessor performance is measured in clock speed, usually in Gigahertz (GHz), which measures the number of instructions per second (in billions) the processor can handle. Examples of PC processors are Intel's Core i7 and AMD's Ryzen. Note that desktop and notebook PCs usually have a single CPU, while servers can have two or more.

Chipset, BIOS, Clock and I/O

- **Chipset (North Bridge and South Bridge):** The chipset supports the CPU by performing some related processing functions and by connecting the CPU out to the rest of the system. The north bridge connects the CPU out to main memory and also to the graphics processor (many include a graphics processor on-chip). The south bridge is the gateway to other PC interfaces (besides graphics) and provides access to the hard disk drives and optical drives via the ATA/SATA, as well as I/O, networking, multimedia, and other peripherals via the PCI or PCI-Express bus.
- **USB Controller and USB Host (integrated into south bridge):** The USB host sits in front of the USB port and receives input from and sends output to a peripheral device using USB (Universal Serial Bus). USB has been widely adopted for its ease and speed: devices connect seamlessly and sometimes require no drivers. USB 2.0 can run at speeds up to 480 Mbps.
- **IEEE 1394/Firewire Controller and Host (not shown):** An alternative standard to USB is IEEE 1394/Firewire. Chipset leader Intel has not pushed this standard; thus most PCs that support Firewire use discrete controllers.
- **BIOS:** BIOS stores boot instructions read by the microprocessor whenever the computer starts up. BIOS is usually implemented with low density NOR flash memory.
- **Clock:** Motherboard clocks provide timing and synchronization functions, keeping all systems in the PC working together. Examples include clock families from IDT, Cypress, and Realtek.
- **I/O Controller:** The I/O controller interfaces the south bridge with the PC's primary simple input and output devices, including the keyboard, mouse, printer and serial devices.
- **PC Card Controller** (not shown, only used in notebooks): This controller is a specialty I/O device that provides the interface with PCMCIA cards. Controllers are only found in notebook PCs. Note that notebooks that support PCI-Express cards do not need a discrete controller, since this is integrated into the south bridge.

Memory

- **DRAM (Dynamic Random Access Memory):** PCs use DRAM to store program data in use by the CPU at any given time. DRAM is blank when the computer boots up and goes blank again when the computer is shut down. DRAM is implemented on modules that can be inserted and removed from the PC motherboard. DRAM is a commodity measured in MBs, segmented by interface (SDRAM, DDR, DDR-II, DDR-III etc.).
- **Memory Register or AMB I/O (not shown, only used in servers):** Servers typically use higher performance memory modules than PCs. These are known as registered DIMM modules, and include register ICs (similar to a timing IC) directly on the module. Newer fully-buffered DIMMs include a specialty I/O IC known as an AMB I/O IC directly on the module, which manages the interface to the server chipset, which would include AMB I/O logic on chipset. Future memory architectures may use buffers directly down on the server motherboard.
- **NAND Cache (not shown):** PCs also contain a memory cache comprised of NAND flash to store frequently used operating system commands and other files.

Graphics and Audio

- **Graphic Processor (GPU):** The graphics processor is a digital processor optimized to produce 2D and 3D graphics images. In desktops, GPUs are found on separate add-in cards (these cards usually contain DRAM specifically used for graphics); in notebooks they are down on the motherboard. GPUs connect up to the north bridge using PCI-Express x16, though older cards still use the specialized AGP graphics bus. Note that many systems eliminate the GPU by integrating graphics functionality into the chipset. Examples of GPUs include NVIDIA's GeForce and AMD's Radeon family.
- **Sound Processor: (not shown, integrated into south bridge, can also be a discrete add-in board):** The sound processor is a digital processor that produces stereo and surround sound audio processing. Sound processors were originally implemented with add-in boards, and a minority of high-performance PCs implement separate PCI or PCI-Express sound cards with discrete audio processors. However, most PCs utilize the sound processor integrated into the south bridge, with discrete controllers reserved for multimedia desktops and notebooks.

- **Audio Codec:** The audio codec is a specialized digital-to-analog/analog-to-digital converter and amplifier that performs the conversion of audio signals from the digital sound processor to an analog format for use with speakers/headphones (output) and microphones (input).

Networking

Note: Networking ICs, including devices found in PCs and servers, are discussed in further detail later on in the networking section.

- **Analog Modem Codec and Line Driver:** The modem codec encodes digital signals from the PC into analog signals for transmission over the phone line and vice versa. The line driver is the analog chip that interfaces with the RJ-11 phone line. In many PCs, a “soft modem” is used, eliminating the need for the discrete codec; however, a discrete line driver remains.
- **Ethernet Controller:** The Ethernet controller is a single-chip integrated device that interfaces between the PC and the Ethernet local area network. It integrates both the MAC and PHY. The Ethernet MAC encodes data from the PC into Ethernet frames and affixes overhead and vice versa. The PHY (physical layer) interfaces with the RJ-45 Ethernet cable and converts the serial Ethernet stream into parallel digital form for processing and vice versa.
- **Wireless LAN Baseband/MAC and RF/IF/Power Amp:** Many PCs include 802.11 (WiFi) wireless networking functionality, especially notebooks. In most implementations, a baseband/MAC chip encodes data from the PC into 802.11 or Bluetooth format and interfaces with the radio. The RF/IF and power amp performs the radio transmission and reception and interfaces with the antenna. Note that we have displayed the chipset as a 2-chip solution; in reality chipsets can be less integrated and will split up some of the functions (separate power amps, multiple radios, specialized MACs, etc.); others will include the solution on a single chip.
- **Bluetooth SoC (not shown in diagram):** Some PCs include Bluetooth personal area networking technology for connectivity to peripherals and headsets. Bluetooth is usually implemented with a single chip SoC containing baseband, RF, IF, and power amp. Though they could connect to the chipset using PCI/PCIe, most PC Bluetooth uses an internal USB socket.

Power Management

Note: Our block diagram does not show specific power management components; these are used throughout the board and on add-in boards as well as within and around the PC power supply.

- **Power Management:** Power management content in a PC includes both analog ICs and discretes. Key types include power transistors, power diodes, AC/DC regulator ICs, PFC pre-regulator ICs, linear regulator ICs, battery charging and management ICs, hot-swap controllers and MOSFET driver ICs.

PCs and Servers

Computing Semiconductor Forecast

- PC and server ICs declining at a 1% CAGR from 2016 to 2021 from $70.8B to $66.4B.
- Microprocessors are forecasted to decline at a 1% CAGR
- Main memory to decline at a 2% CAGR over the same period as memory increasingly shifts to the cloud
- Analog IC grows at a 1% CAGR as PCs increasingly incorporate sensors and PMICs, offset by ASP pressure
- Chipset, I/O, graphics and audio face integration and price pressure.

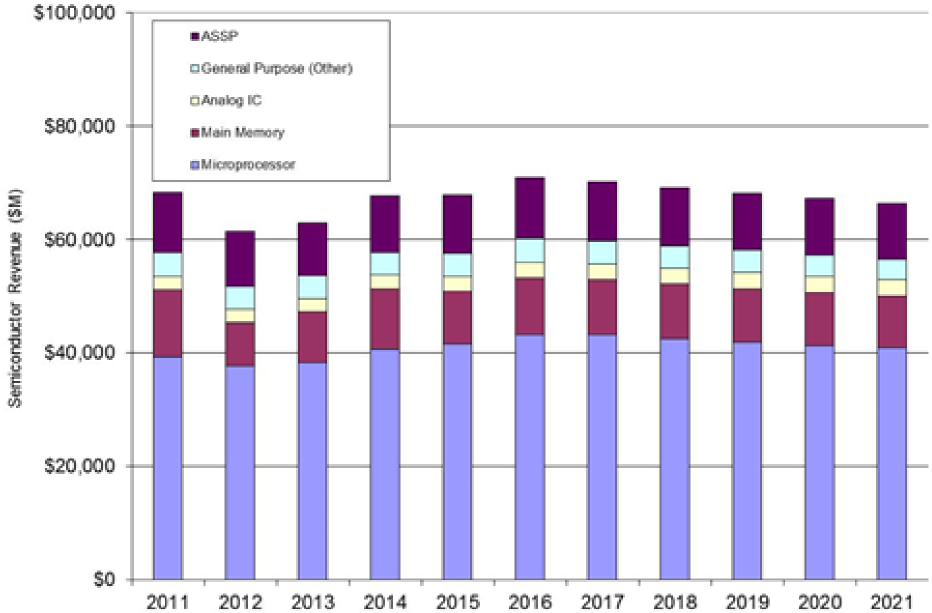

Source: IDC, Gartner, Oppenheimer & Co.

Note: Our entire discussion of PC and server ICs includes chips used in desktop PCs, notebooks, and x86 servers. It excludes chips used in proprietary RISC/CISC servers, as these are typically custom solutions that are difficult to cost within the server bill of materials. It also excludes LCD drivers used in notebook PCs. Also note that our forecast includes all PC and server-related ICs used in full systems, add-in boards, peripherals and the retail after market.

PC and Server Semiconductor Forecast**:** We estimate the computing semiconductor market—which includes desktop, notebook, computing peripherals and x86 server semiconductors—totaled $70.8 billion in 2016, up 4% from 2015. The market began to feel the impact of declining PC units in 2015 after growing 8% in 2014, as well as a shift in the server model away from traditional enterprise computing and towards the cloud.

Looking to 2021, we expect the combination of declining PC units, and the continuing shift to notebooks (which generally have higher component BOMs than desktops) will drive a mostly flattish semiconductor computing TAM. In total, we forecast the market to decline at a 1% CAGR from $70.8B in 2016 to $66.4B in 2021.

Drilling down further, we segment the market by subsystem. The largest device segment in 2016 was microprocessors, at $43.2 billion, roughly 61% of industry revenues (computing and peripherals) by our estimate. As the core processing device of the computing industry, we would normally expect microprocessor revenue to grow slightly above than PC unit growth; PC CPUs have begun integrating graphics, offsetting ASP declines.

Main memory totaled $10.0 billion in 2016. Memory is subject to commodity-like movements in pricing and can see significant annual changes in component costs, driving outsized growth in 2017.

The analog IC and General Purpose (Other) categories are relatively smaller and represent a number of different analog and discreet components within the system. Computing ASSPs represent ~15% of the total market, or $10.7 billion in 2016. We forecast the ASSP portion of the computing market to decline at a 2% CAGR through 2021.

The exhibits below depict our forecasts for the various PC and server semiconductor sub-systems.

PC & Server Semiconductor Subsystem Forecasts

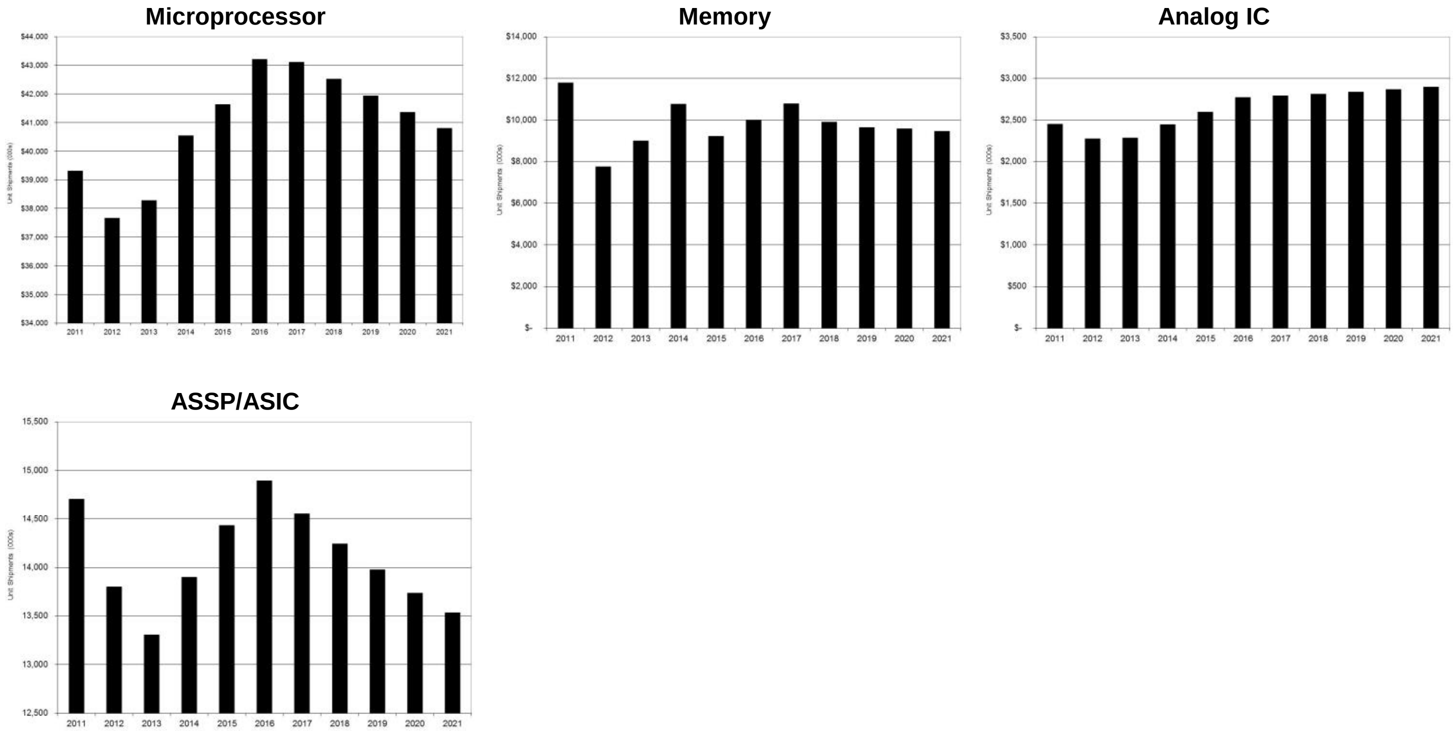

Source: IDC, Gartner, Oppenheimer & Co.

PCs and Servers

PC Semiconductor Market Key Competitors

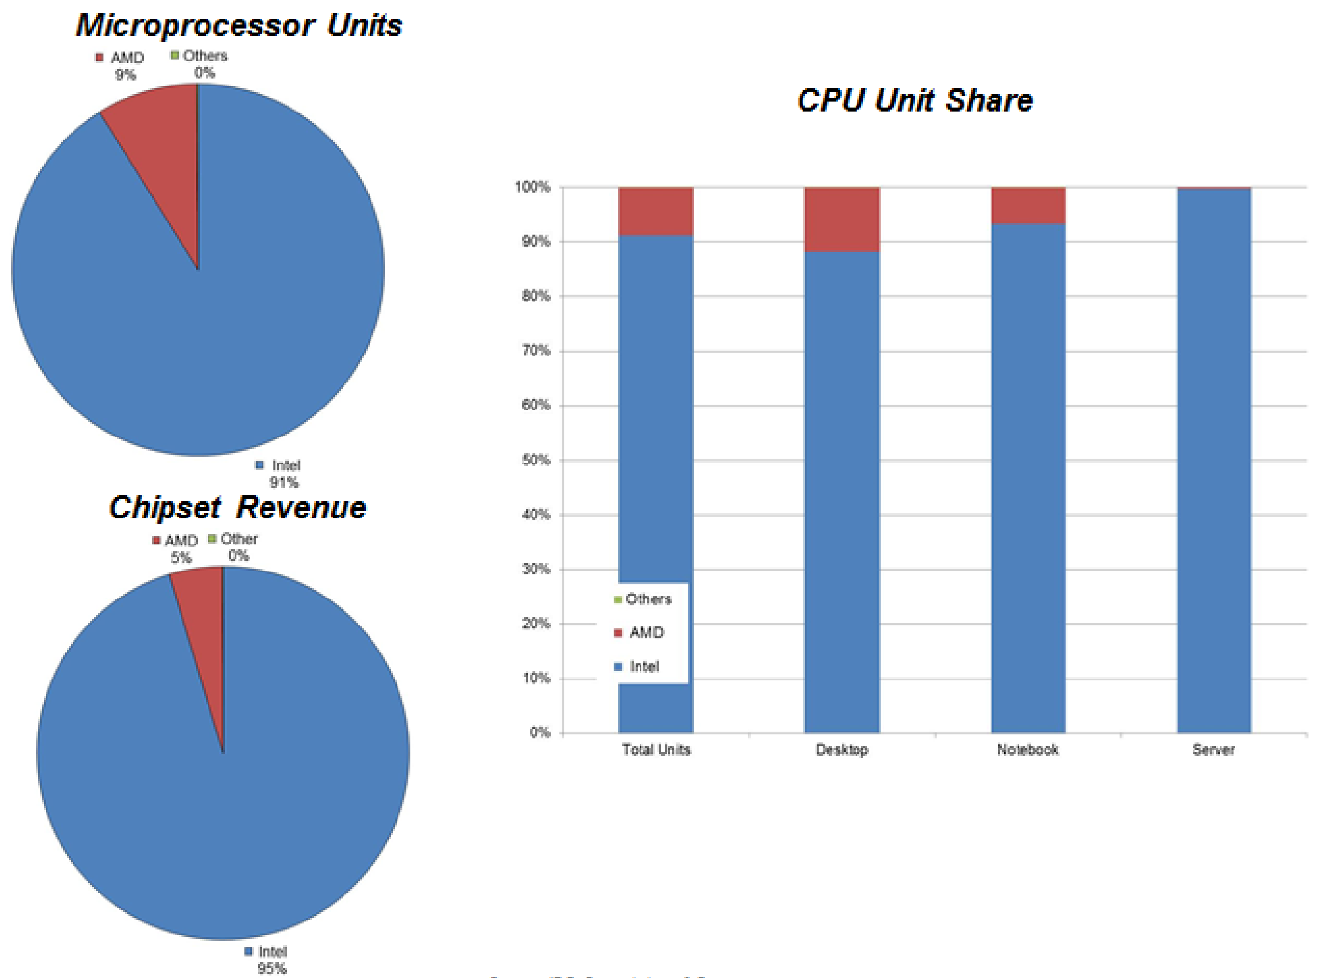

Source: IDC, Oppenheimer & Co.

The PC semiconductor competitive environment is divided firmly along the lines of devices and subsystems. Although many vendors play in several devices within the PC, the standards, design competencies and competitive dynamics in each subsystem are sufficiently different to warrant separate discussion of each subsystem. In the exhibit above, we display 2016 revenue market share data for microprocessors, as well as unit data for desktop, mobile, and server. On the next page, we display discrete graphics controllers and DRAM.

CPUs: The PC CPU market is among the most concentrated in the semiconductor industry. In 2016, Intel held roughly 95% revenue share and 91% unit share. In 2016, AMD held 5% revenue share and 9% unit share. AMD is looking to revitalize its CPU position with its newly released Ryzen PC CPUs and EPYC server CPUs. AMD's presence in server is weakest, and stands at ~1%.

PCs and Servers

PC Semiconductor Market Key Competitors

Graphics

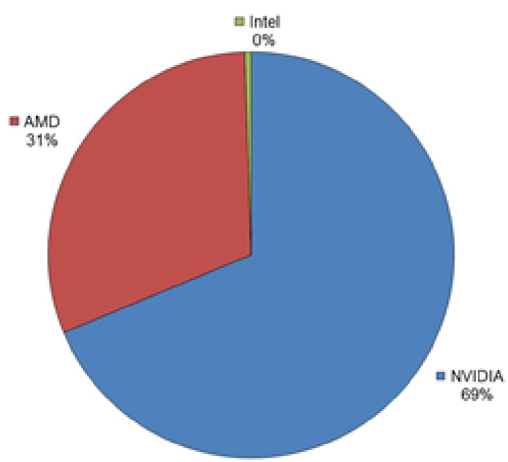

DRAM

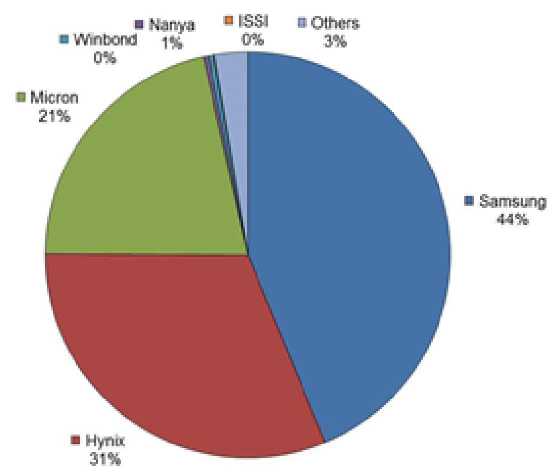

Source: Gartner, Oppenheimer & Co.

Graphics: The discrete graphics market has consolidated over the past few years into the hands of two key players: NVIDIA and AMD. In 2016, NVDA held 69% share of discrete graphics units and this share has been on the rise through 2015. nVidia's revenue share is likely higher than its unit share as it has consistently won high-end share in recent years. AMD is seeking to regain share in the high end with its newly release Vega GPUs.

Memory: The top vendor of DRAM for PCs in 2016 was Samsung, which held 44%. Hynix was No. 2 with 31%, followed by Micron 21%. Note that this data differs slightly from the numbers for DRAM as a whole, which includes non-PC DRAM (generally lower density and lower speed).

Hard Disk Drives

Hard disk drives (HDDs) are still closely tied to the PC market, but drives have dramatically "outgrown" the 1:1 ratio with PCs as enterprise storage networks, the cloud, consumer devices, and personal storage are providing new opportunities for mass storage.

In enterprise, terabytes are growing at a 50%+ CAGR. Networked systems are enabling the separation of processing and storage, allowing for easy, inexpensive capacity additions as corporate data needs grow. Cloud storage is an accelerator to this trend.

Personal storage is a fast-growing application, as increased familiarity with digital content is driving demand for solutions to manage that content.

Hard disk drives (HDDs) are the primary data storage devices used in data processing applications. The drives, usually integrated within larger host devices, use a set of rigid disks to magnetically store bits of data in extremely large quantities at very high speeds. Most PCs ship with at least one hard disk drive, which stores the computer's operating system and software applications and the user's files. They are fast, cheap, carry very large capacities, and conform to industry-standard form factors. They have therefore become an integral part of PCs and servers. However, in recent years, as NAND pricing has become more affordable, SSDs have been cannibalizing HDDs.

Though still closely tied to the PC market, hard disk drives have dramatically "outgrown" the 1:1 ratio with PCs. Enterprise storage, cloud storage, and personal storage are providing new opportunities for mass storage. On the enterprise side, IT managers can supplement their individual PC storage with enterprise storage systems to store and manage common-use applications and files. In these systems, arrays of drives are incorporated into special arrays or networks designed specifically for data storage traffic, which can be accessed by servers on the data network. This dramatically increases storage capacity and performance and also allows the IT manager the option to manage processing and storage independently. Terabytes of enterprise storage continue to grow in excess of 50% per year. The transformational migration to cloud data centers is serving to accelerate this enterprise trend.

Set-top boxes (STBs) and game consoles are proving to be effective usage models for HDDs as these markets require capacity points that cannot be served adequately by flash; other smaller opportunities exist as well. A related driver is personal storage, as consumer familiarity with digital content drives demand for storage solutions to manage and protect that content.

Hard Disk Drives

Hard disk drives are self-contained devices that use rigid disks coated with a magnetic material to store bits of data. They are available in standard form factors using standard interfaces, differentiated by storage capacity, speed, power consumption, and reliability.

Desktop drives are designed for the highest capacity at the lowest cost.

Mobile drives are smaller and consume less power.

Enterprise drives focus on speed, performance and durability.

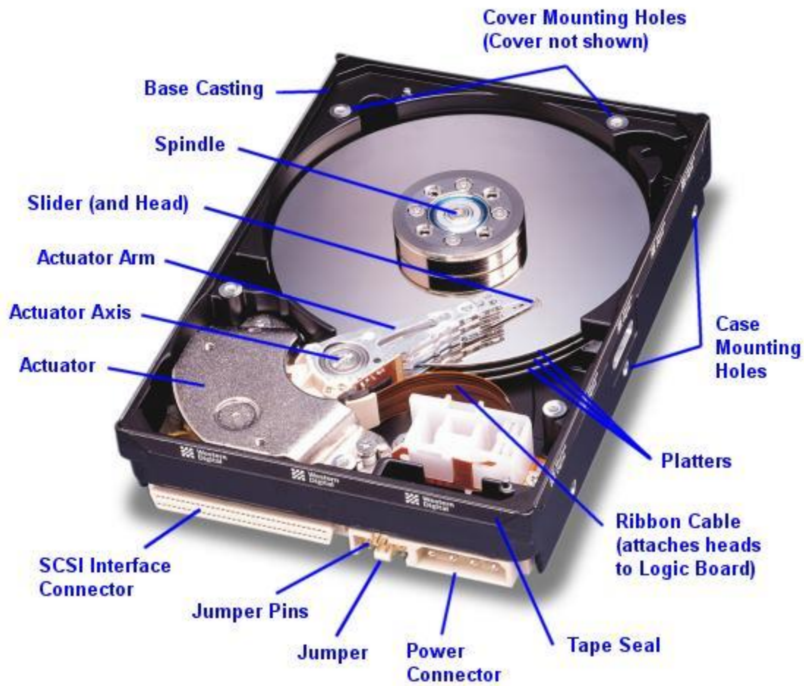

Source: Western Digital

Hard disk drives are self-contained devices that use one or more rigid disks coated with a magnetic material to store bits of data. They are available in standard form factors using standard interfaces, but are differentiated by storage capacity, speed, power consumption and reliability.

A hard disk drive has three basic subassemblies:

- **A set of platters—which are flat, round, rigid disks—used as the storage media.** Each platter is made up of a substrate that is coated with a special media material designed to store data in the form of magnetic patterns. The platters can be double-sided and are mounted on a spindle, which is attached to a spindle motor that spins the platters at high speed.
- **A mechanical device known as an actuator** sits in the corner of the drive and controls a set of arms that reach out over the surface of the individual platters. Mounted on the end of each arm is a slider and a specialized electromagnetic read/write device called a "head." The read/write heads, which are in essence tiny electromagnets that perform the conversion from electrical signals to magnetic signals, are the interface between the magnetic physical media on which data are stored and the other electronic components on the hard disk, recording data to the disk and reading from it. Hard disk drives need one head for each platter surface (i.e., two are needed for a double-sided platter).
- **A printed circuit board, which contains all the logic needed to run the drive system.** Devices on the board include a preamp, which amplifies signals between the head and the drive logic; a read channel, which performs read and write functions; and a hard disk controller and microprocessor, which control the system and implements the interface protocol. The disk drive also has some memory and motor control ICs for the spindle motor and the actuator.

Disk Drive Classes

While the technologies used are very similar across platforms, hard disk drives fall into three basic categories:

- **Desktop.** Desktop-class hard disk drives are used in desktop PCs and low-end workstations and servers. They have standard physical dimensions: 4" wide by 5.75" deep by 1" high (though some drives are less than 1" high) and are generally referred to as "3.5-inch drives." Most desktop drives use the ATA/SATA interface and have one or two platters.
- **Mobile.** Mobile-class hard drives are used in notebook PCs and small form factor desktops. They are smaller than desktop drives; the standard size has dimensions of 2.75" wide, 3.95" deep and 0.37" high and is generally referred to as a "2.5-inch drive." Some are even smaller and are often referred to as "micro-drives"; these drives are found in consumer devices such as portable media players. Most mobile drives use the ATA/SATA interface and have one or two platters; they typically consume less power than other drives.
- **Enterprise.** Enterprise-class hard drives are used in workstations, servers, networked storage systems and high-end desktop PCs. Most have similar physical dimensions to desktop drives, though some high-capacity drives can be greater than 1" in height while others adopt the 2.5-inch mobile form factor. Most enterprise drives use the SCSI interface, though drives designed for network storage often use Fibre Channel. Newer drives are implementing SAS (serial-attached-SCSI), the follow-on to SCSI. Most enterprise drives have more than one platter.

Note that the vast majority of hard disk drives are internal drives, meaning that they are designed to be implemented inside a PC or enterprise system or other "host" device. These generally conform to the specifications indicated above. External drives are sold as standalone units, which are fully enclosed and incorporate additional interfaces such as USB, FireWire, or Ethernet.

Drive Storage Capacity

Storage capacity in a hard disk drive is measured in gigabytes—colloquially known as a billion bytes but actually 1,073,741,824 bytes (computer specifications usually use binary numbers; therefore, a kilobyte is 1024 bytes, a megabyte is 1024 kilobytes and a gigabyte is 1024 megabytes, which gets us to 1,073,741,824 bytes). Note that a byte is eight bits, which are the 1s and 0s of the digital world.

Desktop class drives generally have the highest capacity points over the greatest range, catering to corporate PCs, consumer PCs, and the retail market. Most DVRs and personal storage solutions use desktop class drives as well. Mobile drive capacity points are slightly lower and have a tighter range; these are primarily used in notebook PCs but are also found in portable media players, game consoles, and other consumer electronics. Enterprise drives typically trail desktop drives in capacity, and also have a tighter range.

A drive's capacity is a function of: 1) the amount of bytes that can be stored on each platter (this in turn is determined by the size of the platter and the platter's areal density); and 2) the number of platters in the drive. In general, drive vendors seek to maximize the amount of bytes that can be stored on each platter (the first factor) and to minimize the number of platters in the drive (the second factor). However, in the end the final platter and density implementation will depend on a number of tradeoffs that include engineering constraints, cost dynamics and marketing requirements.

Drive Speed

The speed of a hard disk drive is usually measured on two fronts: transfer speed—which refers to the rate at which data is transferred from the disk to the host system and vice versa; and positioning speed—the rate at which the drive system can hop around the disk in search of the desired data. While many components within the drive affect these metrics, one of the most significant is the strength of the spindle—so much so that spindle speed is usually quoted right after drive capacity when describing a drive and is also the most commonly used specification used to classify different drives by performance. Spindle speed describes the rate at which the platters in the disk drive turn, expressed in revolutions per minute (RPM). The faster a platter spins, the faster the transfer rate when reading or writing as well as the positioning when searching for the correct track and sector.

Spindle speed differs for different classes of drives. In the desktop market, most drives spin at 7,200 RPM, though some higher-end desktop drives spin at 10,000 RPM. Mobile drives spin at speeds ranging from 5,400 RPM to 7,200 RPM (the notebook market requires drives of varying power consumption to meet the different platform requirements). Enterprise drives also vary in speed: from 7,200 RPM at the low end to 15,000 RPM at the very high end, with the most common being 10,000 RPM.

Interface

Hard disk drives use one of several standard interfaces to send data between the drive and the host system. Up until the last few years, the most common interface was IDE (Integrated Development Environment), also known as ATA (Advanced Technology Attachment), which was used in most desktop and mobile drives. ATA drives are configured as master or slave using jumper pins on the drive itself, allowing for two drives to connect into each ATA interface controller on the host. They use wide 80-pin ribbon cables (18" in length) and connect directly to PC motherboards (most PC chipsets include one or more ATA controllers embedded in the south bridge). The data transfer rate for ATA increased over time to 133 Mbps.

The evolution of the IDE interface is known as Serial ATA (SATA), which has quickly displaced ATA. SATA changes the physical architecture from master-slave to point-to-point, so each drive is connected to its own interface. SATA also changes the cable interface to a narrow four-wire cable and extends the length to one meter. Transfer rate was increased to 150 Mbps in the first implementation, with 300 Mbps coming to market now and a roadmap planned to 600 Mbps. SATA is also able to use lower voltage power supplies.

While ATA/SATA has quickly become the most popular PC disk drive interface due to its lower cost, SCSI (pronounced "scuzzy") remains an important interface in the enterprise market as well as in high-end desktops. The SCSI interface uses a "daisy chain" architecture allowing for up to 15 peripherals to be connected to a single interface controller (each peripheral has a second port to connect to the next device). SCSI has also consistently maintained a transfer rate advantage over ATA, with most SCSI drives in use today carrying a transfer rate of 320 MBps. Further, SCSI was the first interface to support RAID (Redundant Array of Inexpensive Disks) configurations. For all of these reasons, SCSI dominates the enterprise hard disk drive market, though the advent of Serial ATA may change the dynamics somewhat. Note that as with ATA, SCSI drive cables connect to a SCSI controller on the PC or server motherboard; unlike ATA, however, chipset vendors have not integrated this functionality into the PC chipset. Servers and SCSI-enabled desktops therefore require a separate controller chip or add-in board to support SCSI drives.

The follow-on to SCSI—known as Serial Attached SCSI (SAS)—is ramping. SAS expands the number of supported devices to 128, uses smaller cable connections (like Serial ATA) and provide universal interconnect compatibility with SATA while retaining the reliability, performance and manageability advantages of SCSI.

Many high end drives incorporate support for Fibre Channel, a high-speed transport technology used in storage area networks. Fibre Channel serializes SCSI commands into Fibre Channel frames, which can be sent at high speed (1.0625 Gbps, 2.125 Gbps, 4.25 Gbps, or even 10.625 Gbps) either point-to-point or through a switched network of Fibre Channel switches. Most Fibre Channel networks, along with their associated arrays of hard disk drives, are set up independent of the network of servers used for data processing, allowing any server to access data stored on any hard disk drive without a direct connection. See the discussion of enterprise storage in the networking section of this report for a more in-depth discussion of Fibre Channel.

Hard Disk Drives

HDD Unit Forecast

- We expect total HDD units to decline at a 7% CAGR, from 387 million units in 2016 to 273 million units in 2021.
- Mobile class drives should decline at 14%, as units will increasingly shift over to solid state drives
- Desktop drives should decline at 7%, from 82 million units in 2016 to 58 million units in 2021
- Enterprise class drives should remain flat as overall enterprise demand grows, offsetting slower flash adoption

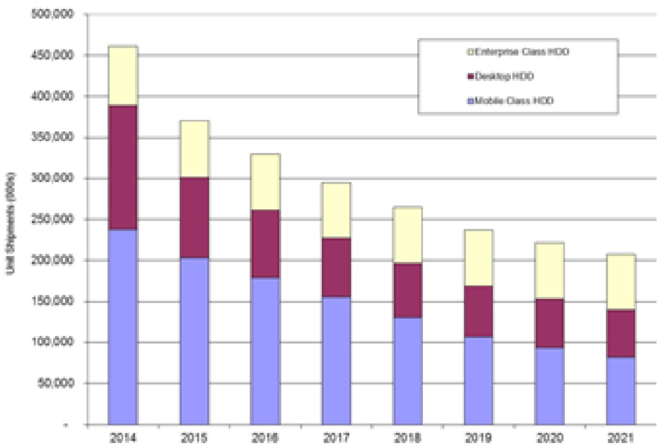

HDD Suppliers

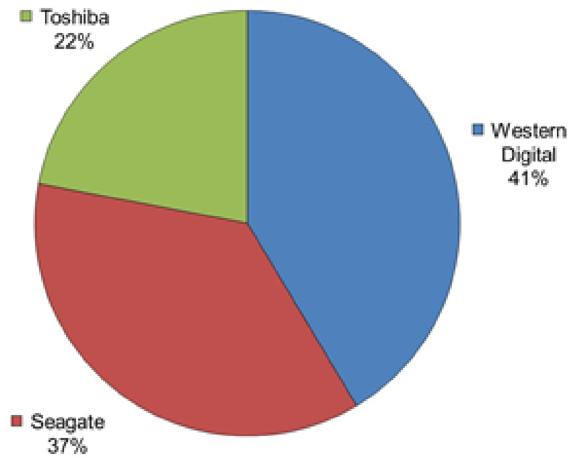

Desktop	***Mobile***	***Enterprise***
Seagate	Western Digital	Seagate
Toshiba	Toshiba	Western Digital
Western Digital	Seagate	Toshiba

Source: IDC, Oppenheimer & Co.

Hard Disk Drive Unit Forecast: Though it originated as a computer storage subsystem, the market for hard disk drives is significantly larger in unit terms than that of PCs. In addition to the unit shipments going into PCs the market includes drives sold into mass storage arrays to support servers and data centers, SAN and NAS systems (two types of networked storage systems), personal storage solutions, and emerging applications in the consumer, cloud, automotive and industrial markets. There is also a significant retail market for add-on HDDs, both internal and external. In total, we estimate the market for hard disk drives totaled 387 million units in 2016. We model a 7% annual decline through 2021, to 273 million units. The market has declined every year since 2010, with the largest declines of 10% and 8% in 2015 and 2016, respectively.

We expect mobile class HDDs to decline at a 14% CAGR, over the same period, driven largely by lower units, and an increasing tendency to design SSDs into notebooks where SSD's smaller footprint and speed are vital. We model desktop units to decline 7% annually to 58 million in 2021 and enterprise units to remain roughly flat at ~68 million.

Hard Disk Drive Competitors: The disk drive industry has gone through considerable change over the past two years. Market share has consolidated into the power of three: Seagate, Western Digital and Toshiba. In 2011, Seagate announced the acquisition of Samsung's HDD business and Western Digital announced the acquisition of Hitachi's HDD business. The result has been consolidated power between the two, better pricing leverage for the industry and a higher margin profile for the HDD competitors.

Source: IDC, Oppenheimer & Co.

***Hard Disk Drive Suppliers*:** As previously discussed, the HDD supply chain has gone through considerable consolidation over the past several years, leaving the industry with only three suppliers. Seagate acquired Samsung's HDD business and Western Digital acquired Hitachi. This left Western Digital with a 41% share, Seagate with 37%, and Toshiba with 22%.

The exhibits below display 2016 unit share data for desktop, mobile, and enterprise hard disk drives.

2016 Desktop HDD Unit Share

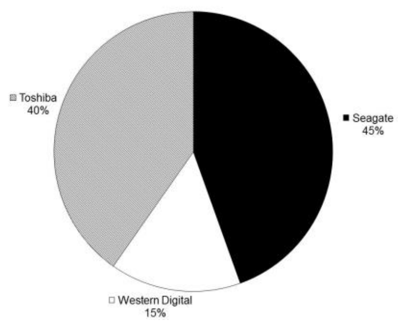

2016 Mobile HDD Unit Share

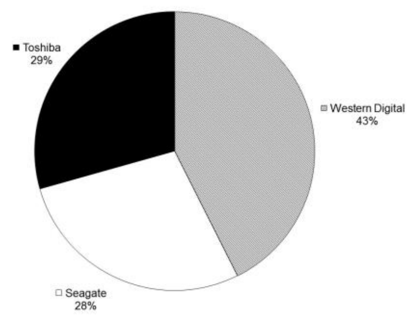

2016 Enterprise HDD Unit Share

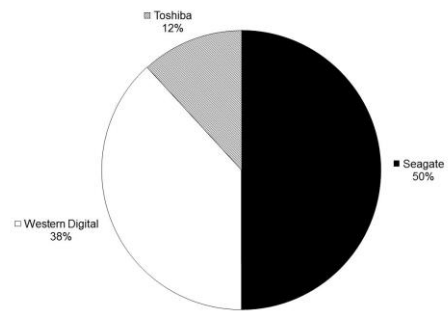

Source: IDC, Oppenheimer & Co.

Hard Disk Drives

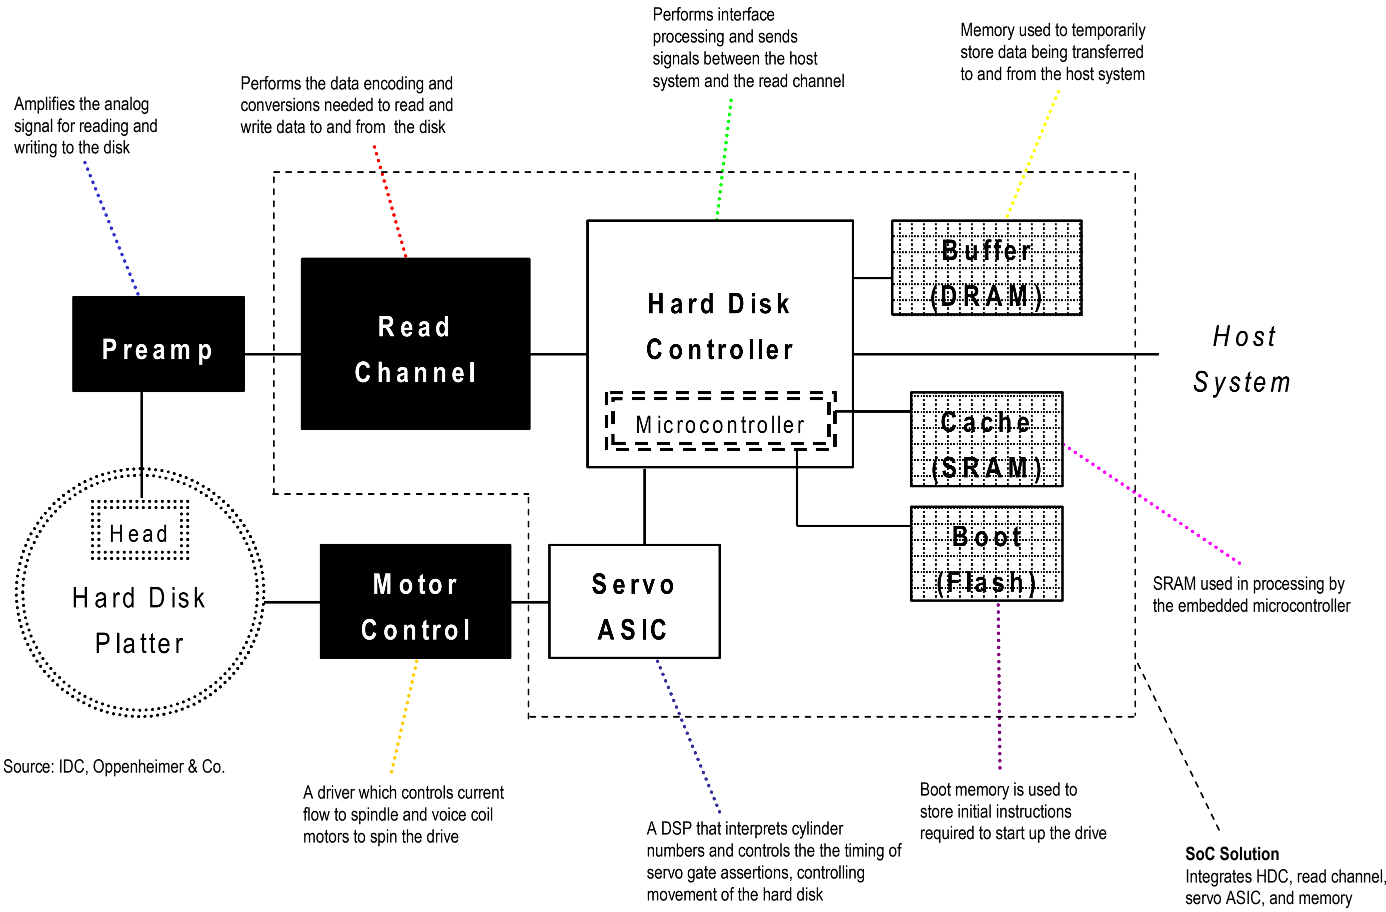

Source: IDC, Oppenheimer & Co.

Note: the diagram above shows a hard disk drive implemented with discrete ICs for each separate function of the system. In practice, most HDDs are implemented with an SoC solution which combines read channel, HDC, servo ASIC, and memory in a single IC.

Semiconductors in hard disk drives perform two basic functions: control of the movement of the hard disk spindle and actuator, and read/write of data to and from the hard disk.

Hard Disk Read/Write

- **Preamp:** The preamp amplifies the analog signal from the read/write heads on the platter and sends the data to the read channel.
- **Read Channel:** The read channel performs the data encoding and conversions needed to read and write the data to and from the hard disk.
- **Hard Disk Controller (HDC):** This controller is the interface processor that sends instructions from the host system to and from the servo ASIC and the read channel. It also determines when to raise and lower read, write and servo gates and is responsible for defect management and error correction. Hard disk controllers usually include a microcontroller to run the entire HDD system.
- **Buffer Memory (DRAM):** Buffer memory temporarily stores the data being transferred to and from the host system.
- **Cache Memory (SRAM):** Cache memory is used to store processor instruction code once the drive is up and running.
- **Code/Boot Memory (Flash):** Boot memory is used to store initial instructions required to start the drive.

- **System-On-a-Chip (SoC)** (not pictured in block diagram above): As mentioned above, most HDDs incorporate an SoC solution that includes most of the read/write functions on a single IC. Most SoCs today incorporate the read channel, HDC, servo ASIC and code and cache memory. Some SoCs also integrate buffer memory, removing the need for a separate DRAM module. The exhibits below display block diagrams without and with SoCs.

Hard Disk Movement Control

- **Motor Control:** The motor control IC controls the amount of current to the spindle motors and the voice coil motor (used to move the head stack), controlling drive commutation and speed. Motor controllers sold as ASSP solutions are specifically designed for hard disk drives.
- **Servo ASIC:** The servo ASIC is a DSP that interprets cylinder numbers and controls timing of servo gate assertions and events during the assertion periods. This digital device is most often integrated into the hard disk controller or SoC.

Hard Disk Drive Without SoC Integration

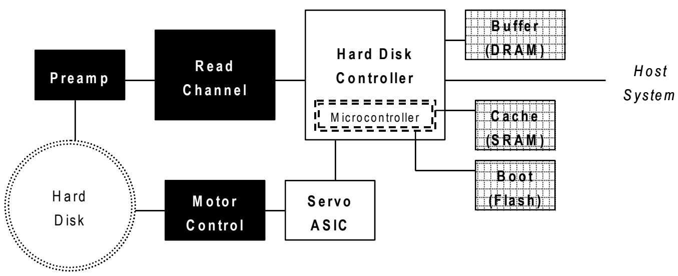

Hard Disk Drive With SoC Integration

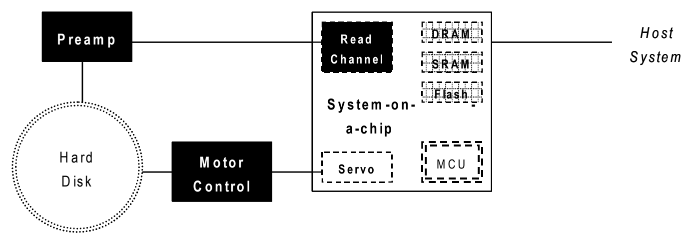

Source: IDC, Oppenheimer & Co.

Solid State Drives

Solid State Drive Forecast

•We forecast total solid state drive (SSD) units will grow at a 15% CAGR, from 157 million units in 2016 to 313 million in 2021.
•Client notebook PCs are the largest market today, representing 37% of units. We forecast 16% unit growth through 2021
•Client desktop PCs will likely drive the most rapid growth (30% annually) through 2021, with consumer electronics also at 21%.

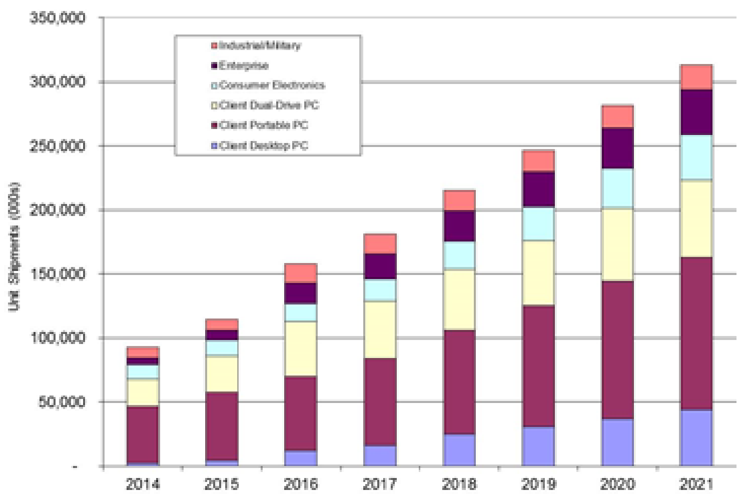

Source: IDC, Oppenheimer & Co.

SSD Competitors

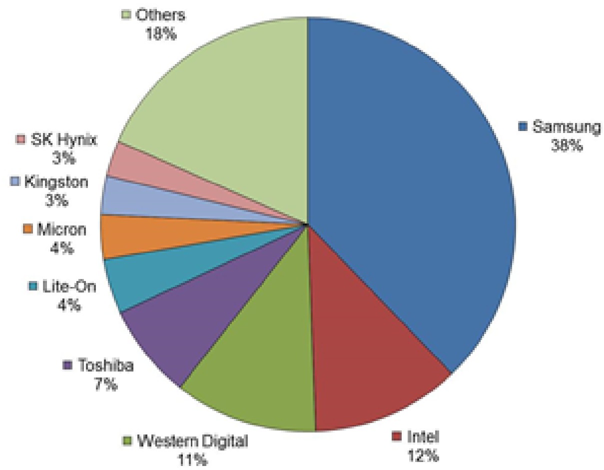

Source: IDC, company reports, Oppenheimer & Co.

Solid state drives have become economically feasible and are poised for continued rapid unit growth over the next five years. Relative to traditional HDDs, solid state drives (SSDs) result primarily in superior data transfer speeds. SSDs are more resistant to physical shock, have lower/faster access time and far less latency. Most SSDs use NAND-based flash memory and prices are today roughly ~7x a traditional HDD. However, falling NAND prices and rising volume have led to increased adoption of SSDs, initially in ultra-portable notebooks and Ultrabooks where fast memory access is paramount. As PCs offload storage to the cloud, less on-device storage is needed. This dynamic, combined with falling NAND prices has made SSDs a viable, if not preferred alternative to HDDs given their superior performance. Gartner estimates most PCs seldom require more than 200GB of storage and that a 256GB SSD will cost ~$40 in 2018. The lowest option it presents for an HDD is a 500GB drive costing ~$34; however we believe the SSD's use case and performance make it the preferred option in most cases.

We forecast total SSD units to grow at a 15% CAGR over the next five years, from 157 million units in 2016 to 313 million in 2021. SSDs grew 38% and 23% Y/Y in 2016 and 2015, respectively. The two largest components of this growth will come from consumer electronics and consumer desktops (in terms of percent). We believe hybrid solutions (combination of SSD and HDD) in desktops will take some share over the next several years before SSDs become entirely economically feasible. While price is a primary inhibitor to adoption today (2-4x HDD for comparable storage), we expect SSD prices to continue to decline.

***SSD Competitors*:** Samsung (38% market share) and Intel (12%) dominate the SSD market today. Western Digital, through its acquisition of SanDisk, holds 11%, while Toshiba owns a 7% share. Lite-On, Micron, Kingston, and SK Hynix are other notable suppliers.

Networking

The computing applications we have discussed generally allow users to manipulate data stored locally within a single device, such as a PC, a server, or a peripheral. We now move on to networking applications, which enable the transfer of data between these devices.

A variety of technologies are used to connect individual computing clients, with a key factor being distance. This section deals primarily with equipment used to connect networking clients in a single physical location, e.g., a small business, a single corporate branch office, or a home network environment.

Equipment used to send data over longer distances (e.g., modems) typically utilizes public telecom and cable infrastructure or corporate wide area networks and are discussed later in the section on telecom and datacom. Cellular networks are also discussed separately.

The computing applications we have discussed generally allow users to manipulate data stored locally within a single device, such as a PC, a server, or a peripheral. We now move on to networking applications, which enable the transfer of data between these devices.

A variety of technologies are used to connect individual computing clients, with a key factor being distance. This section deals primarily with equipment used to connect PCs, servers and storage in a single physical location—e.g., a small business, a single corporate office, or a home network environment. Given the close ties to computing, networking equipment is usually designed with the PC in mind. Metrics like speed, ease of install and configurability and low maintenance cost generally take precedence over factors such as ensuring quality of service or multiple classes of service support (a situation that is reversed in telecom networks).

Equipment used to send data over longer distances typically utilizes public telecom and cable infrastructure or corporate wide area networks and is discussed separately (in the section on telecom and datacom). This equipment is designed to leverage the existing voice or video infrastructure to send data rather than act as a standalone network, though some standalone wide area data networks do exist. The economics are also different, as telecom and datacom equipment is tied to service provider capital expenditures, while networking is tied more closely to IT spending.

Cellular technologies are also discussed separately in the section on wireless, as the technology and economic factors driving those technologies are worthy of separate attention.

Networking

The OSI Reference Model

The OSI reference model provides a framework for the development of networking technologies and standards.

The model divides networking functions into seven layers, all of which can be designed independent of each other. The upper layers (4-7) deal with user applications, and function independently of the type of physical network used to carry information. The lower layers (1-3) are dedicated to issues relating to the actual transport of information around the network.

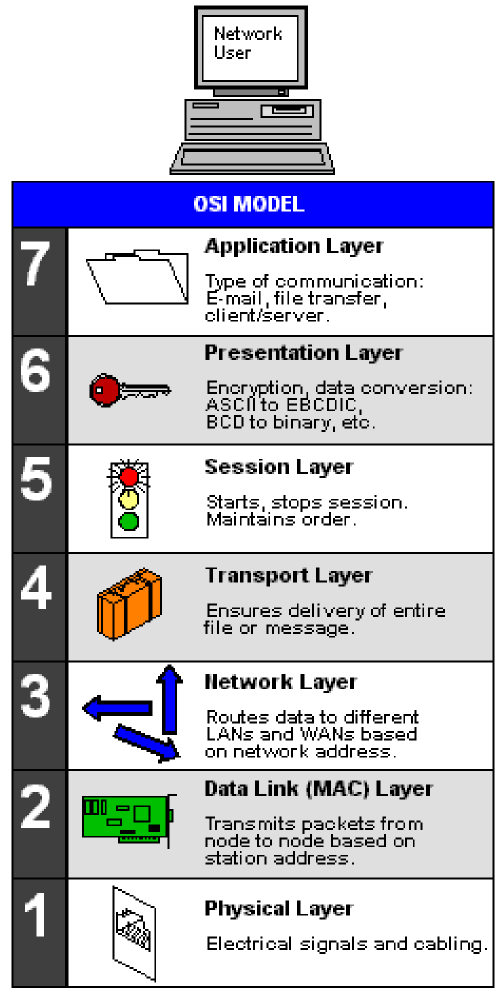

Source: Computer Desktop Encyclopedia, The Computer Language Co.

Networking technologies are designed to share data between computing devices. These devices connect using a wide variety of terminal equipment and operate using disparate communications protocols, transmission media and data transfer rates. Given the disparate technologies, the challenge of internetworking becomes how to connect these networks such that data can be reliably transmitted among them.

In order to meet the challenges of internetworking, network architects began to segment communications tasks into separate and self-contained layers of functionality. Each layer is designed to perform a certain task or range of tasks and is independent of the layers above and below it. This process was formalized in 1984 by the International Organization for Standardization (ISO) with the publication of a seven-layer reference model. The reference model governs how information from a user application in one computer is transmitted across a network or inter-network to a user application in another computer. In this way, it provides a framework for the coordination of standards development, allowing existing and evolving standards to be set within a common framework. The OSI model has therefore become the primary architectural model for internetworking communications.

The layers of the OSI reference model can be grouped into two categories: upper layers (4-7) and lower layers (1-3). The upper layers deal with actual user software applications (like Microsoft Outlook or Internet Explorer), and function independently of the type of physical network used to carry information. The upper layers are typically implemented in software running on the networked device. The lower layers are dedicated to data transport issues and can be implemented in either software or hardware.

In order to understand the function of each layer of the model, we can look at the process of sending an E-mail from a PC. A computer is a layer 7 device, meaning that it can operate from the highest layer (which is where applications are run) to the lowest layer (the PC can put signals directly onto a network cable plugged into a port). At the application layer (7), Outlook (or another E-mail client) deals with the composition of the E-mail, including formatting, attachments, and security features. At the presentation layer (6), the E-mail is converted into a standard format compatible with all E-mail programs. At the session layer (5), a channel is opened between the PC and the network for message transmission. At the transport layer (4), the operating system passes the E-mail to the computer's networking subsystem. At the network layer (3), the system makes sure it has a clear path to the rest of the data network. At the link layer (2), a network address for the E-mail is attached to the message. Finally, at the physical layer (1), the PC sends out the actual signal over the network cable. In receiving an E-mail, the process is reversed, though the functions performed are the same.

The OSI Reference Model: Layer 1-7 Summary

The Physical Layer (1)

This layer defines the electrical, mechanical and functional specifications for transmitting a data signal over a network link. The physical layer is also responsible for establishing, maintaining and terminating the physical signal between network devices. Physical layer protocols typically define parameters such as voltage level, shape or template of transmission signal, data transfer rates, maximum transmission distances and physical connectors.

The Data Link Layer (2)

The data link layer controls the transfer of data across a physical network link. As such, data link layer protocols define the characteristics of network links including network topology, physical addressing, frame formats, frame sequencing, error notification and flow control. The IEEE (Institute of Electrical and Electronics Engineers) further subdivides the link layer into two sub-layers: logical link control (LLC) and media access control (MAC). The LLC sub-layer manages communications between devices separated by a physical link, while the MAC sub-layer manages access to the physical network medium.

The Network Layer (3)

The network layer is responsible for establishing, maintaining and clearing a network-wide connection between two transport layer entities. The network layer principally provides routing and other functions that enable data to be transmitted over disparate physical links. The network layer provides both connection-oriented and connectionless service to layers above it. Most network layer protocols are routing protocols. Note that network layer routing functions are based on logical addressing rather than physical addressing.

The Transport Layer (4)

This layer functions as the interface between the application-oriented upper layers and their underlying network-dependent layers, providing transparent yet reliable inter-network data transport services to the higher layers. Effectively, the transport layer hides the detailed operation of the underlying network from the upper layers. Transport layer functions include flow control, multiplexing, virtual circuit management and error correction. Flow control manages the rate of data transmission between two entities such that the transmitting devices do not send data at a rate greater than what the receiving entity can process. Multiplexing allows several communications channels to be carried over the same physical link without any data loss. Virtual circuit management involves the creation, maintenance and termination of virtual circuits (a virtual circuit is a lower speed channel carried over a higher speed physical link). Finally, error checking involves mechanisms for the identification and correction of transmission errors.

The Session Layer (5)

This layer is responsible for establishing, managing and terminating a communication channel between presentation layer entities through a series of communications requests and responses. In addition, the session layer determines whether communication will be half-duplex (two-way transmission, but only one direction at a time) or full-duplex (simultaneous two-way transmission), maintains synchronization between communicating entities, and performs exception reporting.

The Presentation Layer (6)

This layer is concerned with the syntax of data during transfer, i.e., it ensures that data is formatted using a common interface for user applications. To achieve network interoperability, a number of transfer syntaxes have been defined. The presentation layer negotiates and selects the appropriate transfer syntax for each communication so that data transmitted may be understood by both systems. If the transfer syntax differs from the application syntax, the presentation layer performs the necessary conversion. Some common transfer syntaxes include the ASCII and EBCDIC for text and data, motion picture experts group (MPEG) and QuickTime for video and graphics interchange format (GIF), joint photographic experts group (JPEG) and tagged image file format (TIFF) for graphics images. The presentation layer also performs data compression/decompression and encryption/decryption if required.

The Application Layer (7)

This layer interacts with and provides services directly to user software applications. It deals directly with the communications component within an application (e.g., the "Inbox" in Microsoft Outlook) and provides the interface to the services of the network. Functionality provided by this layer typically includes identification of communication partners, determination of network resource availability, synchronization, file transfer, message interchange services such as E-mail, agreement of privacy (encryption) standards and identification of constraints on data syntax.

Networking Equipment, Components, and the OSI Model

Though it all seems technical and theoretical, it is important to understand the layers of the OSI model in order to identify the intended functionality of the various types of networking equipment. For instance, a layer 1 device would deal only with actual physical signaling.

An example of a layer 1 device would be a repeater (a device used to boost a signal over a length of cable), which simply receives a signal (at the physical layer [1]) and spits it back out.

Layer 2 deals with addressing. An example of a layer 2 device would be a switch, which receives traffic (at the physical layer [1]), checks the address of the traffic (at the data link layer [2]), and sends the traffic out again (at the physical layer [1]) over the proper channel to reach the recipient.

Layer 3 deals with routing. An example of a layer 3 device would be a router, which receives traffic (at the physical layer [1]), checks the address of the traffic (at the data link layer [2]), determines the best route for the data to take to its destination (at the network layer [3]) and then re-addresses and sends out the traffic (at the data link [2] and physical [1] layers).

We won't carry the example all the way up to layer 7, but it is clear why we called a networked PC a layer 7 device. It can do all the functions of a switch, router, or other more complicated devices because of its robust and flexible operating system and processing hardware. In fact, layer 4-7 network devices are often designed as specialized servers designed to run networking software.

We can also carry the analysis forward to the component world (though we are skipping ahead a bit). An Ethernet PHY is a physical layer (1) device, as it deals with reception and transmission of the signal over copper wire or optical fiber. Every piece of equipment that physically connects to an Ethernet network must therefore contain an Ethernet PHY, from a repeater (a layer 1 device) up to a load balancing device or a networked PC (a layer 4-7 device).

An Ethernet MAC is a data link layer (2) device, as it applies and removes an address to a data packet. An Ethernet MAC would be found in every piece of Ethernet equipment that operates at layer 2 and above, such as an Ethernet switch or NIC card all the way up to the load balancer.

A packet forwarding engine is a network layer (3) device, as it looks up addresses in a route table and forwards packets on to their destinations. A packet forwarding engine would be found in Ethernet equipment that operates at layer 3 or above, from routers up to our load balancer.

As mentioned above, layer 4-7 devices are generally implemented in software running on a specialized networking OS. For high-performance equipment, OEMs often use a combination of microprocessors, specialized ASICs, ASSPs, and FPGAs to perform these tasks.

Networking

Communications Flow in the OSI Reference Model

For two PCs to communicate, data must travel down the layer stack to the physical layer so it can be passed between nodes in the network. At each node, the packet is processed and forwarded; how far up the layer stack it is processed depends on the function served by the node.

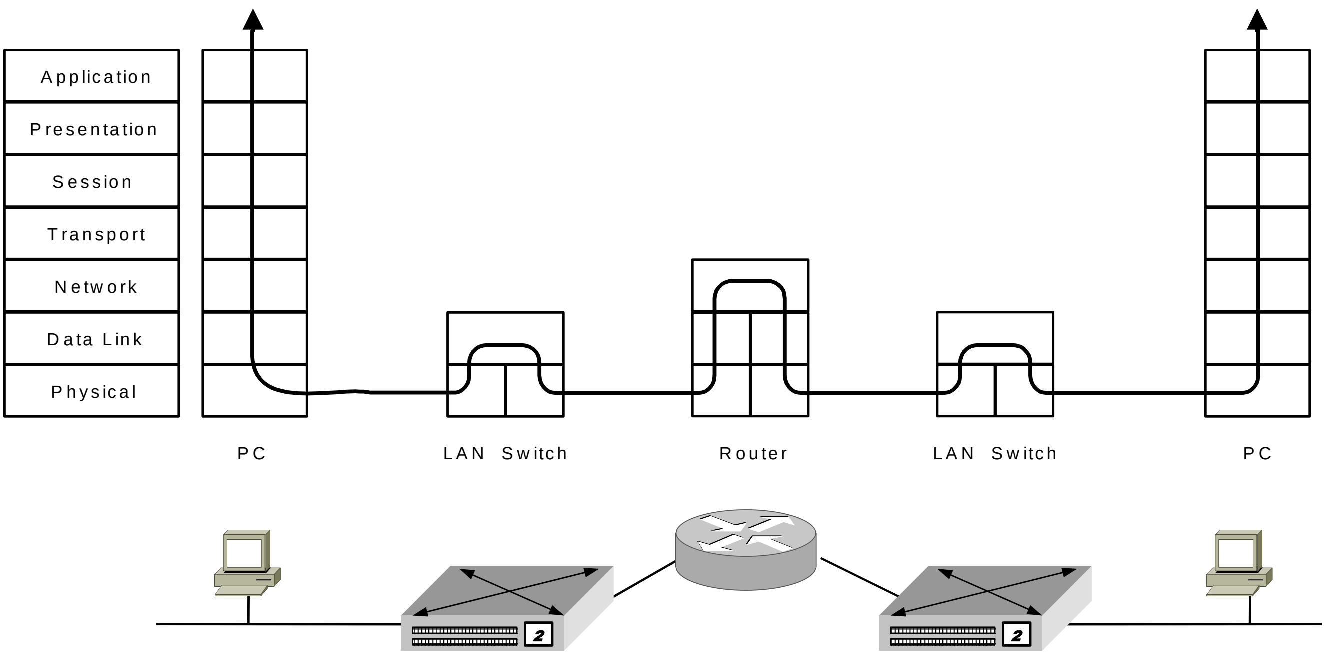

Source: Oppenheimer & Co.

We now look at the OSI reference model in action, and track how it governs communication between networked devices. For data to be transferred from a user application program (e.g., Outlook or Internet Explorer) on one computer (or other connected computing device) to a user application program on another computer, it must pass through each of the OSI layers on both the transmitting and receiving computers. This is because user applications operate at layer 7 of the OSI model, while the port used to transmit the data operates at layer 1.

To get the data from layer 7 to layer 1 on the transmitting device, control is passed from one layer to the next, starting at the application layer, down through the middle layers until it reaches the physical layer. The data then travels as layer 1 traffic over the physical network link to the receiving device and back up the layer hierarchy.

If more than one physical link separates communicating devices, for instance if the devices are on different floors in an office building, the information contained within the transmitted signal will be processed by each intermediate node sitting on the network between those two devices. These nodes can be hubs, switches, routers, servers, or anything else connected to the network. How far up the OSI layer stack the information will be processed depends on the function served by the node. For instance, LAN switches tend to process information at layer 2, meaning that the information will travel up the layer stack from layer 1 to layer 2, will be processed by the switch at layer 2 and then will go back down to layer 1 within the switch before it is sent on its way. Routers process data at layer 3; the path is similar except that the information will travel as high as layer 3 before being sent back down. More complex devices operate at layers 4-7, as they need to look at the type of data being transmitted in order to perform their functions.

Networking

Encapsulation in the OSI Reference Model

Each layer of the OSI model can communicate directly with peer devices at the same layer. In order to communicate with layers above and below, control information in the form of headers and trailers is used. Higher layer headers and trailers are treated as payload data for lower layers. This technique is known as encapsulation.

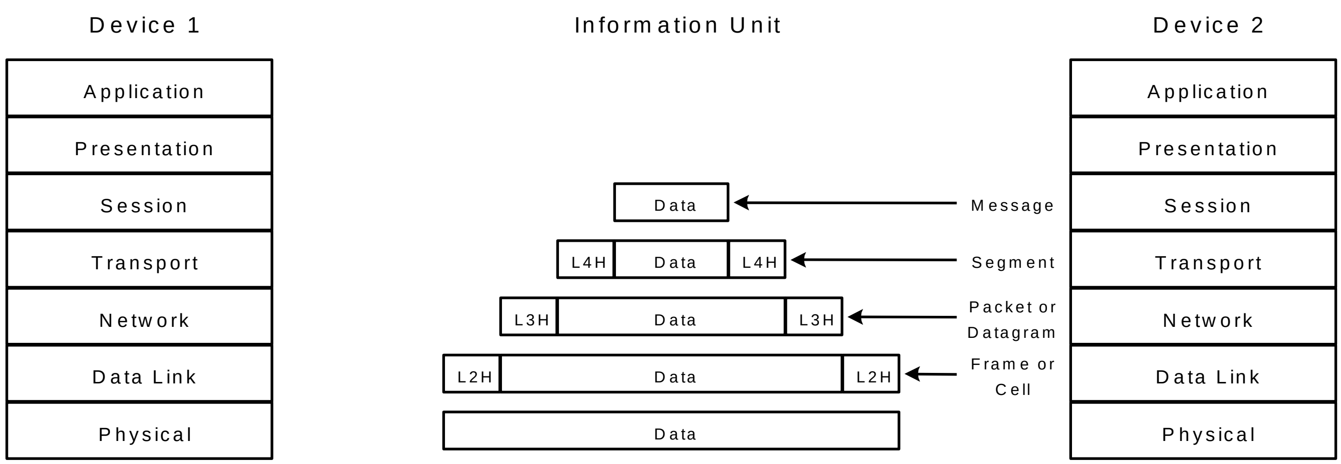

L 4 H	=	Layer 4 header	L 3 H	=	Layer 3 header	L 2 H	=	Layer 2 header
L 4 T	=	Layer 4 trailer	L 3 T	=	Layer 3 trailer	L 2 T	=	Layer 2 trailer

Source: Oppenheimer & Co.

Each layer in the OSI model can communicate with only three other places in the network: the layer directly above it within the device, the layer directly below it in the device and its peer layer in other networked devices.

The latter process, communication with peer devices, is straightforward: at any given layer, networked devices use the same protocols and are designed to interact together. This is dictated by the standards that govern the networking protocol. In order to communicate with other layers, however, control information in the form of headers and trailers prepended or appended to the payload data stream is used. Each layer uses the information in the header/trailer below it in order to process the received payload data. As information travels up the stack, each OSI layer removes the header attached by its corresponding peer layer and forwards the remaining information to the next higher layer. On the way down, each layer will add its own headers and trailers; these will be treated as payload data by lower layers. This is known as encapsulation.

Some common terminology used in referring to the information units passed between OSI layers is defined below.

- *Message:* An information unit whose source and destination layers lie above the network layer.
- *Segment*: An information unit, including header/trailer, whose source and destination layers are transport layer entities.
- *Packet or Datagram:* An information unit, including header/trailer, whose source and destination layers are network layer entities. Datagrams are used in connectionless services.
- *Frame or Cell:* An information unit, including header/trailer, whose source and destination layers are data link layer entities. A frame is of variable length while a cell has a fixed length (like in an ATM network).

Encapsulation Demonstration

To demonstrate the principles of encapsulation, we will follow an E-mail as it makes its way through the PC, and will complete the process by sending it over the network and opening the E-mail on the recipient's PC. An E-mail application at layer 7 on the transmitting PC generates the E-mail and tells the system to send it over the network (that's Microsoft Outlook telling Windows it has an E-mail to send). The PC initiates a session (Windows accesses the communications software on the Ethernet NIC/LOM connection) and puts headers and trailers on it and sends it to the lower layers. As it travels down the network stack (within the NIC/LOM at this point), each layer adds headers and trailers to the outside of the previous layer until a fully stacked frame or cell is created. The frame or cell is then placed on the network by the physical layer device (the PHY device on the NIC card).

The frame or cell, encapsulated in some physical layer protocol (e.g., Ethernet), is then read by the physical layer and data link layers of all other devices to determine if the frame or cell was intended for it. In the case of switches and routers, the devices will forward the frames out toward other network devices (the distinction being that switches will dump the frame out the port closest to the destination, whereas a router will actually find the best route for the frame and forward it on).

Once it finally reaches its destination, the cell will travel up the layer stack. Each layer (starting with those in the NIC card, then the communications software, then Windows) will strip off the headers and trailers from the layer below and pass the data up until it reaches the application (Outlook again or some other E-mail client).

Networking

Four Key Markets:

Ethernet (LAN)

The most pervasive local area networking protocol, Ethernet is used to connect PCs and other networked devices in a single physical location.

Wireless LAN (802.11)

Wireless LANs connect users in a single location using radio waves, usually used in conjunction with a LAN or broadband Internet connection.

Bluetooth

The most common personal area networking protocol, Bluetooth connects two devices in close proximity wirelessly. Bluetooth is particularly prevalent in the handset market.

Storage

Enterprise storage systems, including direct attached storage (DAS), storage area networks (SAN), and network attached storage (NAS).

This section deals with pure data network technologies used to connect multiple PCs, peripherals, or storage devices in a single physical location. Networking equipment is usually designed around the needs of corporate IT networks, though consumer networking has gained significant traction in the past few years. Key technologies include:

- **Ethernet**: The primary wired data networking technology for computers and servers, Ethernet is a fast, scalable, cost-effective technology used to connect PCs in a corporate or home network.
- **Wireless LAN (802.11)**: Wireless LAN mimics a wired LAN but replaces the copper cable with radio waves—allowing for the networking of devices at distances up to 300 feet without wires. It is usually used in conjunction with a LAN or broadband Internet connection. 802.11 uses unlicensed spectrum, primarily 2.4 GHz and 5 GHz.
- **Bluetooth:** A personal area networking technology, Bluetooth is a fast, cheap, low-power wireless protocol used to connect two devices in close proximity. Bluetooth is particularly prevalent in the handset market.
- **Storage:** Enterprise storage, while primarily implemented via direct-attached hard drives today, is increasingly moving to Fibre Channel-based storage area network and network attached storage systems, mimicking an Ethernet network and providing increased efficiency, manageability, scalability and upgradeability.

Throughout the networking section, we will be quoting a variety of speed grades that can be difficult to keep track of. The chart below displays the major networking technologies along with the sub-protocols of each. We also include telecom/datacom, home networking, and industrial networking technologies for comparison purposes. We only include true networking technologies, however; interface technologies such as USB, SATA, SAS, PCI-Express, RapidIO, HyperTransport, and a whole host of others are not included.

Networking Protocol Speed Chart

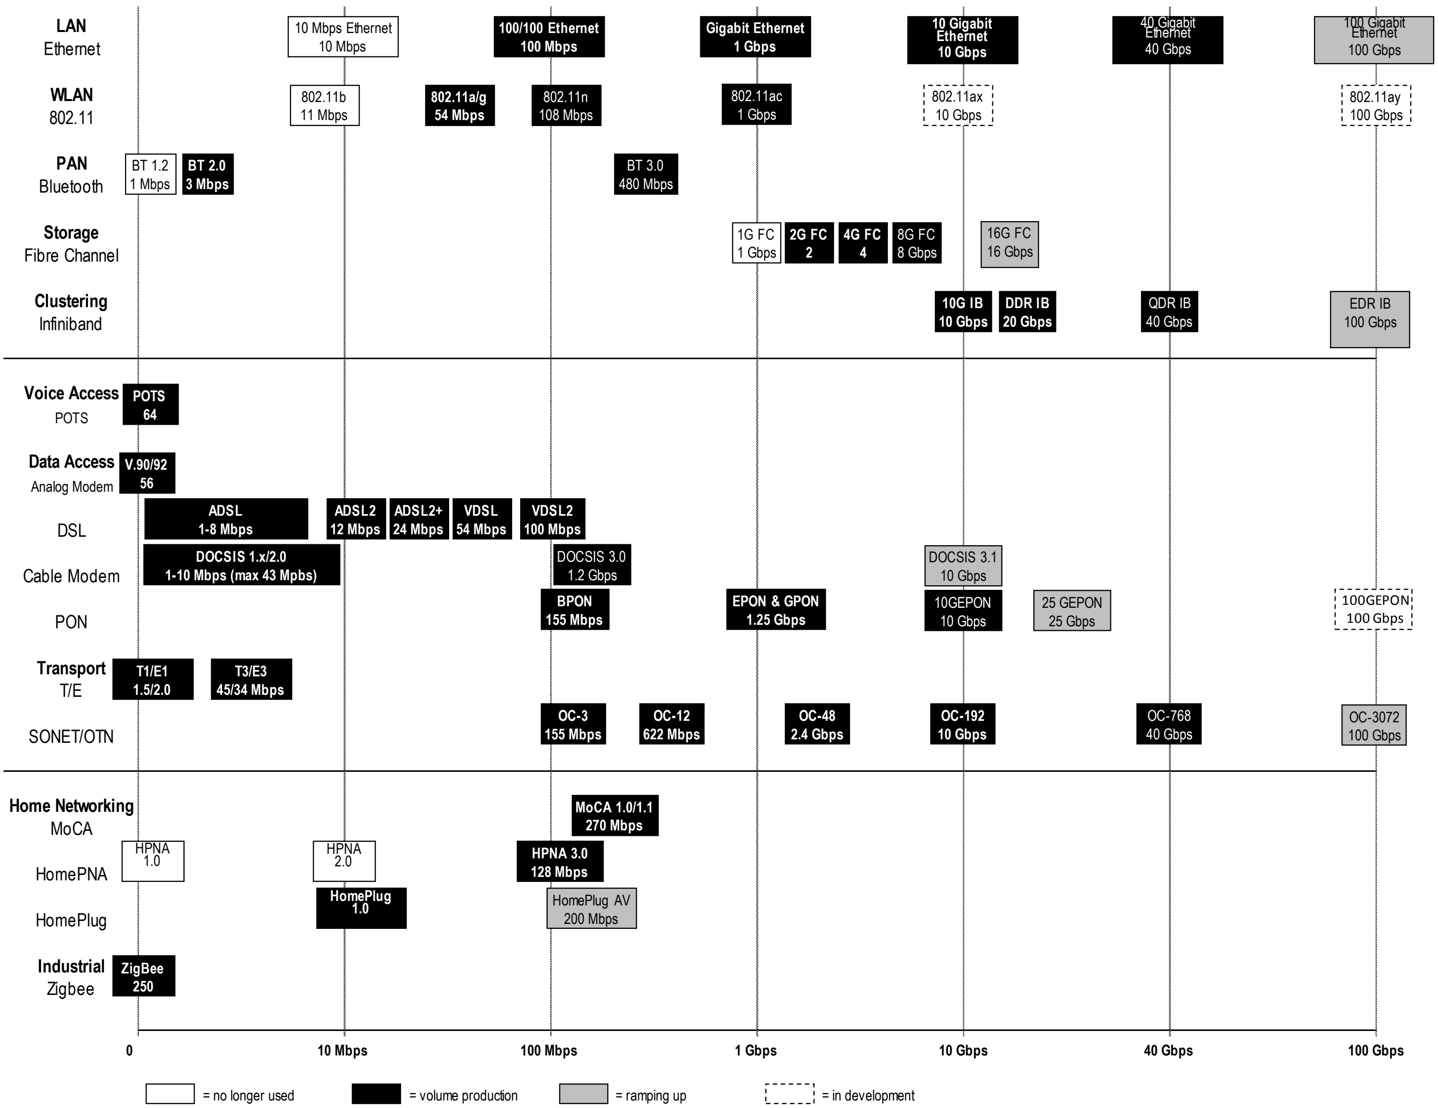

Note: In most cases, speeds quoted correspond to the maximum data rate (sometimes called the "PHY rate)" achievable according to the standard. Most networking protocols have a usable data rate that is meaningfully less than the quoted rate, though the magnitude of this difference can vary significantly depending on the protocol.
Source: Company reports and Oppenheimer & Co.

Ethernet

Ethernet is the most popular protocol used in local area networks. PCs, servers, and printers equipped with Ethernet ports connect to switches that direct traffic around the network. Data packets bound for locations outside the LAN are addressed and forwarded by routers.

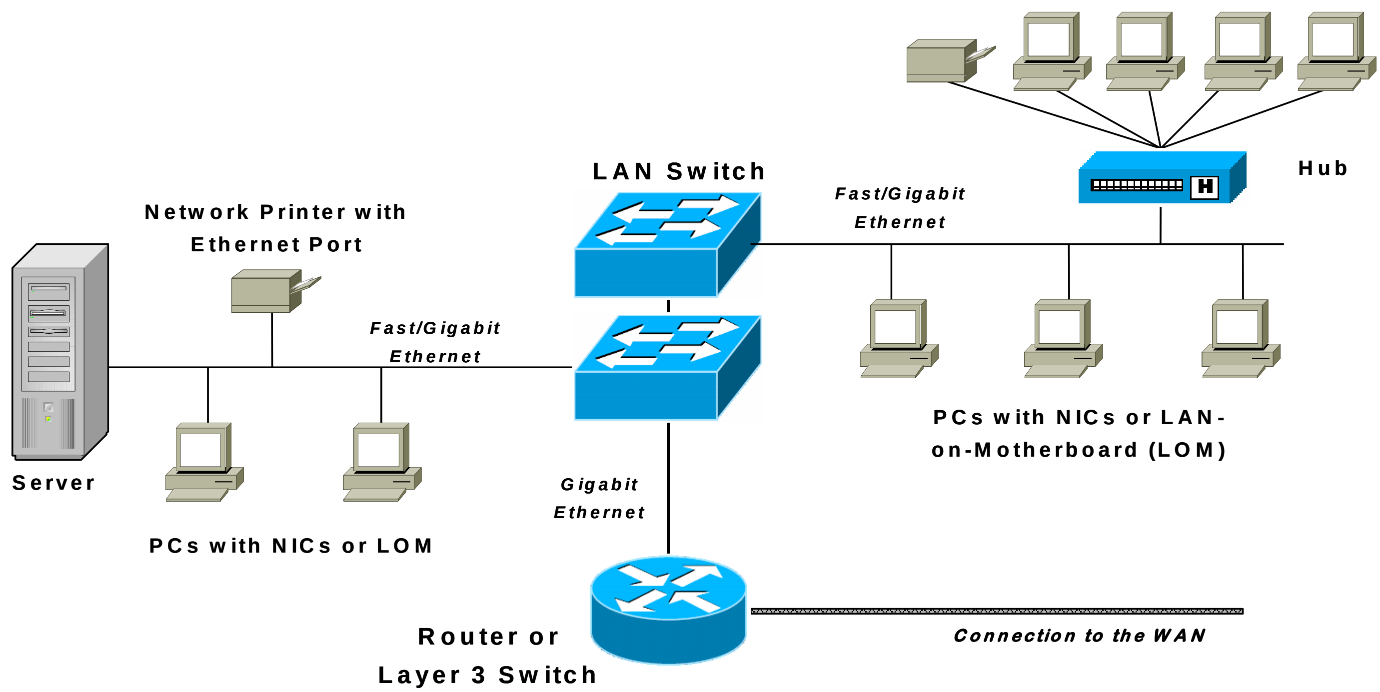

Source: Cisco Systems, Oppenheimer & Co.

Ethernet ***START HERE*** is the most pervasive networking protocol used in local area networks. Ethernet is a packet-switched frame-based broadcast technology, meaning that it sends addressed information out over the network rather than establish a dedicated line between two devices (we discuss packet-switched and circuit-switched networks in more detail in the datacom/telecom section). To get data from one device to another, the transmitting device (e.g., a PC) "broadcasts" the data, along with the recipient's address, over the network. All devices attached to that network (known as "client devices") link read the address; if the packet is bound for that device it extracts the payload; otherwise that packet is ignored. In practice, most Ethernet networks use a switched architecture, meaning that switches are used to forward traffic between clients on the network; this cuts down on redundant traffic and frees up bandwidth on the network.

To connect to an Ethernet network, client devices require a Network Interface Card (NIC). Most PCs today ship with LAN-on-motherboard (LOM) solution, which is essentially a single-chip NIC implementation soldered directly onto the PC motherboard. Each client connects via Category 5 cable up to an Ethernet switch (fiber can also be used). As client devices broadcast Ethernet frames, the switch reads the addresses and determines which of its ports the recipient sits on and broadcasts the packet out on that port, where it is received and opened by the recipient client.

Switches in the LAN are generally connected to one another via uplink ports (an uplink port is the same as a standard Ethernet port but is sometimes a higher speed port) to form a LAN backbone. This LAN backbone allows communication between clients throughout the wider LAN and also interfaces with a router to send and receive data to locations farther away, such as another LAN or the public Internet.

Ethernet

Types of Ethernet Equipment

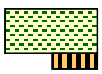

NIC (Network Interface Card) - Add-in card installed in PCs, servers, and other Ethernet clients to connect the device to the network. Most client devices have migrated to LAN-on-motherboard solutions.

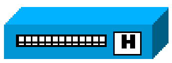

Hub - Simple LAN device that connects multiple clients by forwarding Ethernet traffic to each port without any intelligent switching.

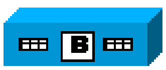

Bridge - Simple LAN interconnect device used to connect two LANs.

Layer 2 Switch - The basic network building block, layer 2 switches connect LAN clients together and switch traffic between them, forwarding Ethernet frames between specific ports based on the network address.

Router - Intelligent packet forwarding device that operates at layer 3 of the OSI stack. Routers are placed strategically within a network and communicate with each other to determine the optimal route through the network for packet forwarding. Routers also perform protocol translation.

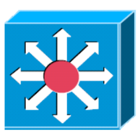

Layer 3 Switch - Adds routing capability to the basic Ethernet switch, eliminating the need for packets to be routed centrally.

Layer 4-7 Switch - Performs packet forwarding and routing based on deeper packet information (as per layers 4-7 of the OSI model), enabling multiple classes of service (CoS) and better quality of service (QoS). These switches can also do load balancing and intelligent provisioning.

Types of Ethernet Equipment

Network Interface Card (NIC): Network interface cards are add-in cards that provide individual clients with an interface to the LAN. NICs operate at the lowest two layers of the OSI reference model. They are responsible for segmenting and framing the user data stream (data link layer functions) and for signal shaping and conditioning (physical layer functions) required to transmit data over the LAN medium. IEEE compliant NICs ship with a predefined, hard-coded MAC address that uniquely identifies the terminal device from all others connected to the LAN.

LAN-on-Motherboard (LOM): Most PCs now ship with a LAN-on-Motherboard solution. These single-chip solutions essentially capture the entire functionality of a NIC card on a single device, which is soldered on the PC motherboard.

Hub: Hubs connect multiple network clients but do not perform any intelligent switching. Basically, a hub simply forwards Ethernet traffic to each port in order to extend the LAN to additional clients. They provide congestion control in LANs and also serve as a point of data aggregation in the LAN.

Bridge: Bridges are simple LAN interconnection devices that connect two bus-based LANs running the same protocol. Like hubs, bridges operate at the lowest two layers of the OSI reference model (the data link layer and the physical layer) and forward traffic from one LAN segment to another.

Switch: Switches connect LAN clients together and are the basic building blocks of a local area network. Like hubs and bridges, switches operate at the data link layer and physical layer of the OSI reference model and provide traffic forwarding based on destination addresses (known as MAC addresses). Switches differ from bridges and hubs, however, by offering intelligent traffic forwarding between multiple input ports and multiple output ports through a central switch fabric. They therefore prevent network congestion on LAN segments that do not include either the originating or receiving device. Furthermore, switches allow multiple simultaneous communications across the fabric on a point-to-point basis.

Router: Routers are intelligent packet forwarding devices that operate at the three network dependent layers (physical, data link and network layers) of the OSI reference model and serve two basic functions: determining optimal routes or paths through the network, and packet forwarding. Routers typically serve to interconnect LAN segments as well as to connect LAN segments to the WAN. As such, routers typically support multiple data networking protocols, such as T1/E1, T3/E3, Ethernet, Token-Ring, FDDI, Frame Relay, ATM and SONET/SDH. Routers, especially high speed routers, perform the packet routing in hardware in order to provide the greatest possible performance.

Unlike switches that view the network on a link-to-link basis, routers view the network on a global basis and have knowledge of neighboring routers. Routers dynamically maintain network topology data, including neighboring routers, route distance, route delay and route congestion levels, and update this information as network conditions change. In addition, routers are self-learning devices. When introduced into a network, a router exchanges information as to network topology with neighboring routers. In this process, the newly introduced router learns the topology of the network from its neighbors, while simultaneously allowing its neighbors to update their view of the network to reflect the existence of the new router. When making packet forwarding decisions, routers depend on this accumulated network intelligence. Routing decisions may be based on such factors as destination address, payload type, packet priority level (contained in the IP header), least-cost route, route delay and route congestion levels.

Layer 3 Switch: As discussed, switching of Ethernet frames between clients on the LAN typically takes place at layer 2 of the OSI reference model. By adding layer 3 functionality to a switch, a hybrid switch/router is created that can forward and switch packets at layer 3, choosing a specific route for the packet to follow through the network to its destination. By making intelligent routing decisions when the packet is switched, the number of network routers may be reduced, thereby limiting the number of router peers with which each router must interact.

Layer 4-7 Switch: Layer 4-7 switching operates at the transport layer (layer 4) and above of the OSI reference model. Layer 4-7 switching goes beyond IP to transport layer protocols such as TCP or UDP to make packet forwarding decisions. Layer 4-7 switches are able to look deep into data packets to extract information regarding the packet content and the application protocols exchanging the data packet. Using this information, together with details about network congestion, delay and loss, the layer 4-7 switch makes packet forwarding decisions. With the ability to differentiate classes of traffic, layer 4-7 switches have the ability to prioritize those flows with more strict requirements, thereby meeting quality of service demands. They also perform load balancing functions, keeping congestion down to manageable levels.

Ethernet

IP Traffic Driving Demand

- Global IP traffic is expected to rise from 96,054 PB/month in 2016 to 278,108 PB/month in 2021, a 24% CAGR.
- The surging number of mobile connected devices and expanding video/OTT capabilities is driving demand for Ethernet infrastructure.
- Switch ports are still growing from both new LAN buildouts and healthy upgrades; networking OEMs continue to pack better performance and more advanced security.

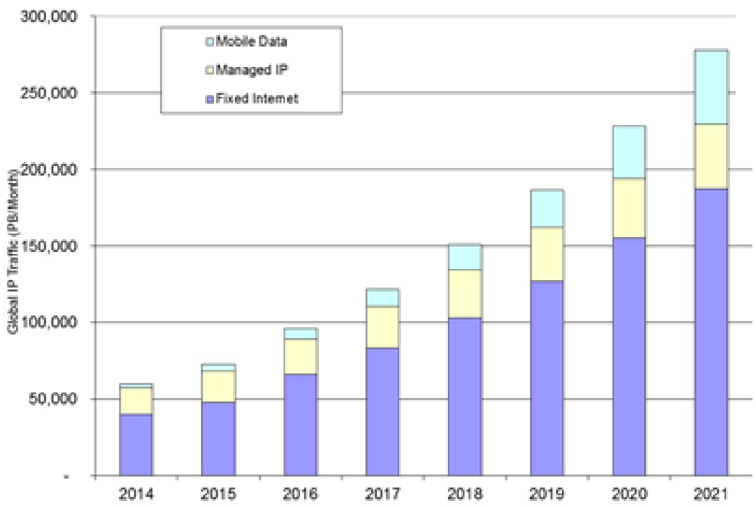

Source: Cisco Visual Networking Index, Oppenheimer & Co.

Ethernet Equipment Suppliers

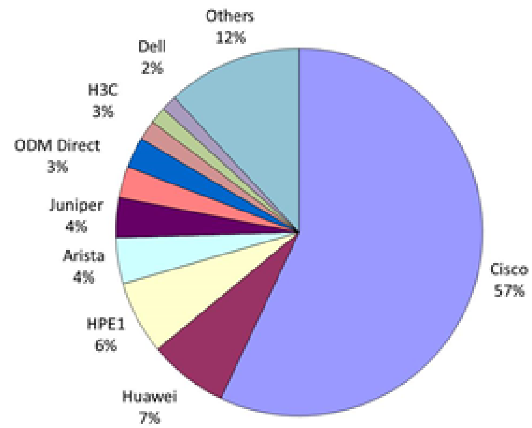

PCs	*Servers*
Lenovo	HP
HP	Dell
Dell	IBM
Acer	Fujitsu/F-S
ASUS	NEC
	Sun

Source: IDC, Oppenheimer & Co.

***IP Traffic is Driving Demand**:* With the surging number of mobile connected devices and subsequent data usage demands, annual global IP traffic is poised more than double again from 2016 to 2021, after growing nearly 5x from 2010 through 2016. This is pushing the bandwidth and efficiency limits of current network infrastructure and Ethernet has emerged as the de facto standard to enable this surge in traffic.

A major opportunity to capitalize on this growth is Ethernet switches, which include simple layer 2 switches (simpler devices that switch and forward packets), layer 3 switches (which add routing functionality) and layer 4 and above switches (switches with higher layer intelligence, for load balancing, QoS, and other management capabilities).

Infrastructure equipment, which includes blade servers, storage networking devices (iSCSI), broadband infrastructure (DSLAMs, CMTS, etc.), and other telecom equipment, is also likely to experience robust growth through 2015.

***Ethernet Equipment Suppliers**:* Cisco was the dominant supplier of Ethernet switches in 2016 with 57% market share, up from 34% in 2012. HP continues to lose share and is down to 6% in 2016, from 19% in 2012. Huawei, Arista and Juniper round out the rest of the top players, with 7%, 4%, and 4%, respectively.

Vendors of non-PC client devices include manufacturers of modems, printers, game consoles, set-top boxes, and a variety of other PC peripherals and consumer electronics devices.

Ethernet

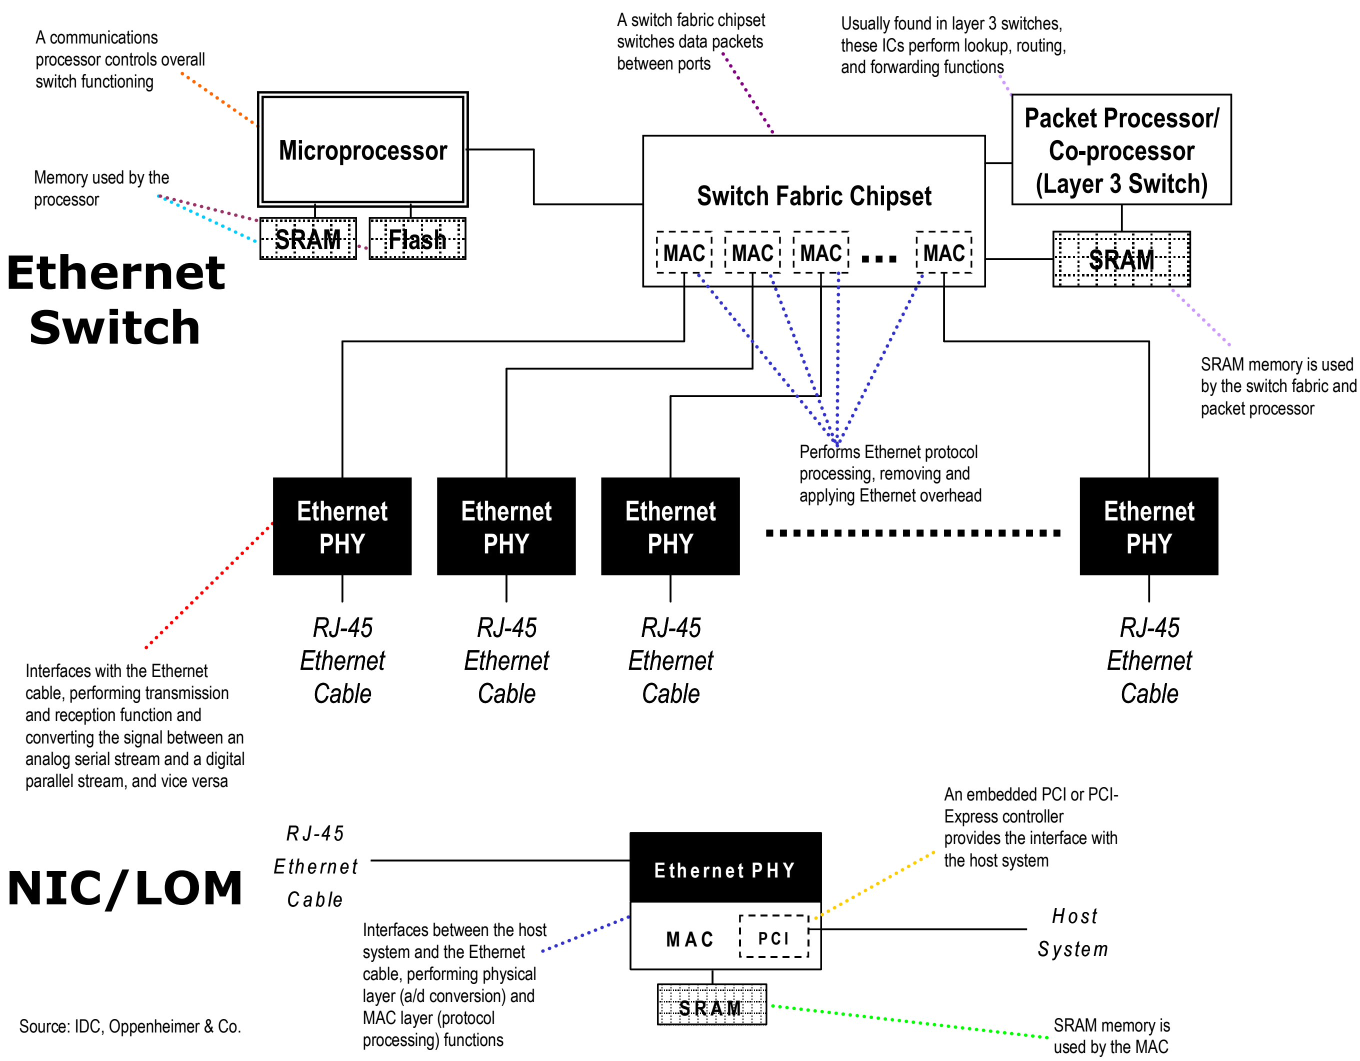

Source: IDC, Oppenheimer & Co.

Ethernet Port *(NIC and Switch)*

- **Ethernet PHY**: Sometimes called a transceiver, the PHY interfaces with the RJ-45 Ethernet cable (or optical fiber) and converts the serial analog stream to a digital parallel stream and vice versa. In a switch, these are implemented as multi-port (quad, octal, etc.) PHYs. In a NIC/LOM, they are part of the Ethernet controller.
- **Ethernet MAC, SRAM**: The MAC, sometimes called simply the "controller," assigns and removes Ethernet headers and overhead. In a switch, these are usually integrated into the switch fabric chipset. In a NIC/LOM, they are part of the Ethernet controller. MACs typically use a small quantity of SRAM memory.

Switching Block (Switch Only)

- **Switch Fabric:** The switch fabric receives packets from the various Ethernet ports and switches them between ports or sends them to the uplink port (if the destination is a port on another switch or the WAN or Internet).
- **Packet Processor/Co-processor** (layer 3 switch only): Found in layer 3 switches, packet processors perform routing and forwarding functions. A packet co-processor (NSE/CAM) performs intelligent lookups.
- **Microprocessor, SRAM, Flash**: This block—which can contain a microprocessor, FPGAs, SRAM and flash—controls the overall system and performs high-end functions.

Ethernet

Ethernet Semiconductor Forecast

- We expect the enterprise networking market to grow from $32.9B in 2016 to $39.3B in 2021, an 4% CAGR.
- Layer 2-3 switches will remain the largest revenue driver for semiconductor suppliers, remaining at ~60% of the market
- We forecast Fibre Channel and Inifiband to decline at ~6% CAGR through 2021

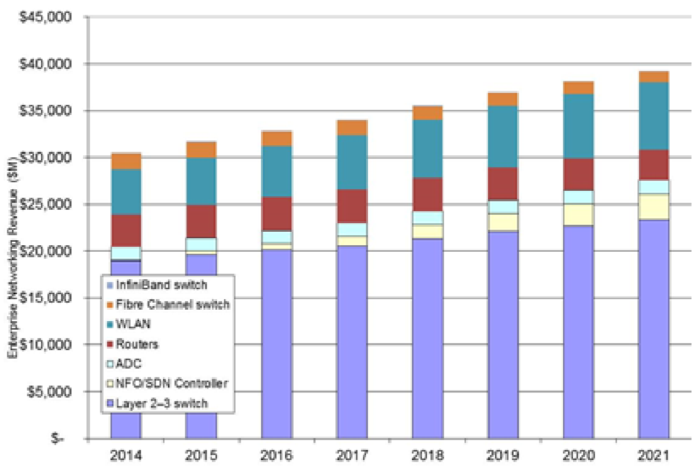

Source: IDC, Oppenheimer & Co.

Cloud/Mobile Comm Semis Competitors

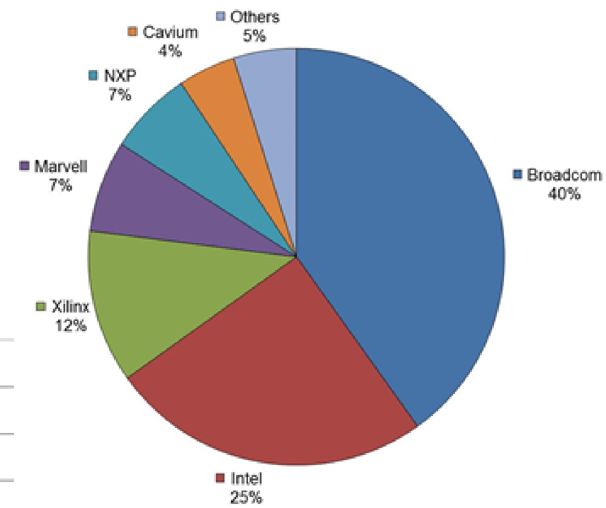

Source: IDC, Oppenheimer & Co.

Semiconductor Revenue in Enterprise Networking Forecast: The market for enterprise networking equipment is expected to grow at an 4% CAGR, rising from $32.9 billion in 2016 to $39.3 billion in 2021. Semiconductor Ethernet suppliers will continue to benefit from several ongoing upgrade cycles, including the transition from SONET/SDH to OTNs. The IEEE 802.3ba 40G/100G Ethernet standards were ratified in 2010, and Ethernet will continue to play an ongoing role in the build out of the cloud and in datacenters – all major data center vendors now support either standalone or integrated Ethernet solutions.

Ethernet Semiconductor Competitors: The top suppliers of cloud/mobile infrastructure semiconductors are Broadcom Intel, Xilinx, NXP, and Marvell.

The No. 1 supplier in 2016 was Broadcom, which held 40% share. Broadcom serves all segments as mentioned above, and is particularly dominant in the server NIC and Gigabit switching segments.

Intel, is No. 2 with 25% share, predominantly serves the PC market with Ethernet controllers organically, but also serves the market through its recently acquired Altera FPGAs. Xilinx's 12% is primarily derived from selling FPGAs to the comm infrastructure market while Marvell supples Ethernet semiconductors and storage controllers. Rounding out the top suppliers are NXP and Cavium.

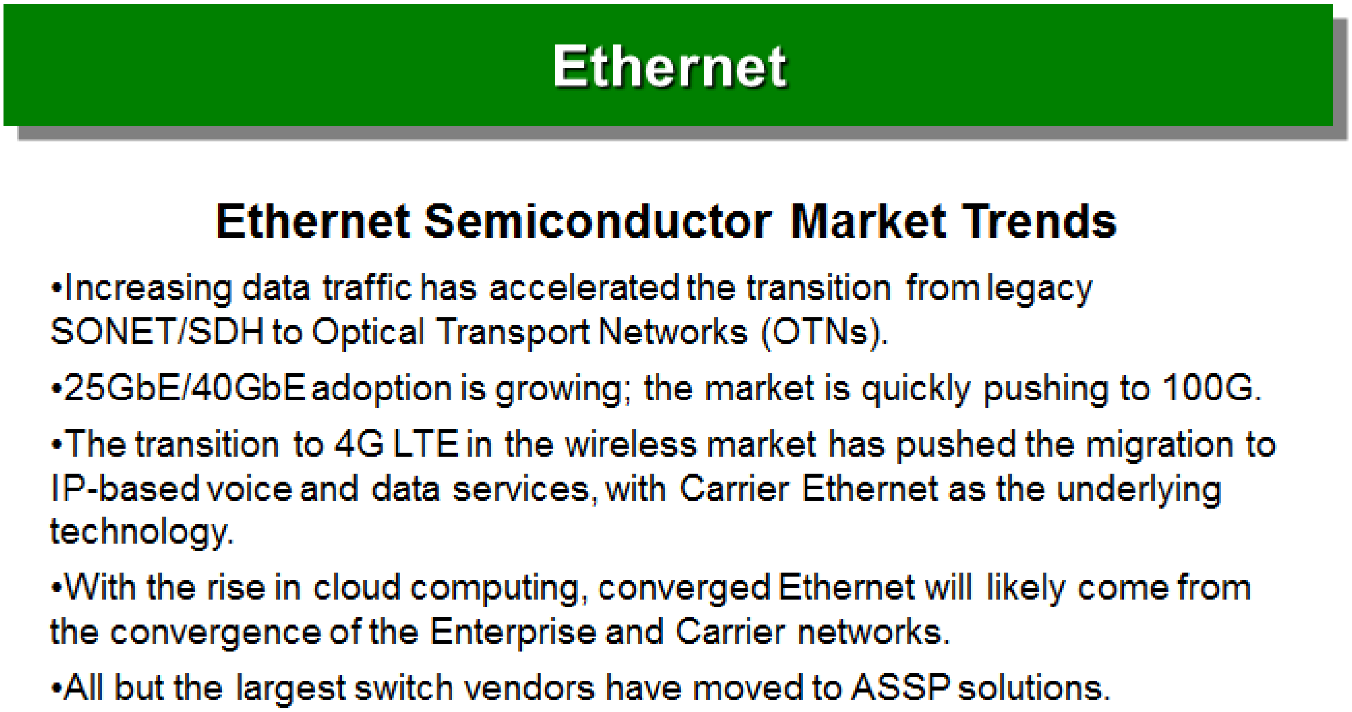

- **Increasing data traffic is driving the transition from SONET/SDH to Optical Transport Networks (OTNs).** Data traffic is undergoing an exponential growth phase, expected to more than double again from 2016 to 2021. Driven by the dearth of mobile connected devices and the new ability for consumers to constantly stream video, share photos etc., carriers are under increasing pressure to implement cost-effective capacity expansion. OTN has emerged as the Layer 1 protocol of choice.
- **40GbE/100GbE adoption is accelerating; the market is quickly bypassing 10GbE.** 10GbE semiconductor revenue declined ~6% in 2016. Growth has decelerated as the demand for data has driven more rapid adoption of 40GbE and 100GbE ports which grew 15% and >5x respectively Y/Y in 2016. We expect the market to continue to transition to 100GbE at an accelerating rate over the next 2-3 years.
- **The transition to 4G LTE in the wireless market has pushed the migration to IP-based voice services, and, Carrier Ethernet as the underlying technology.** As service providers rush to meet growing video and data requirements on the wireless networks, the transition to 4G has occurred faster than any previous network upgrade cycle, and the pending upgrade to 5G only amplifies this trend. LTE is expected to ultimately migrate to using IP as the standard communication protocol, and carrier Ethernet will be the underlying technology standard. We would expect the transition to occur over the next several years.
- **With the rise in cloud computing, converged Ethernet will likely come from the convergence of the Enterprise and Carrier networks.** The explosion of cloud computing has begun to move enterprise applications and data into secure hosted applications which can be accessed from multiple points and form factors. This substantially increases the demand for data centers to build out high bandwidth, high reliability networks that can interconnect this infrastructure. Converged Ethernet will likely result as the carrier and enterprise networks meet at the edge of the cloud, driving further demand for semiconductor solutions.
- **All but the largest switch vendors have moved to ASSP solutions.** As the market has continued to move to Gigabit, 10GbE, and 25GB+, most switch vendors have moved, or are moving to ASSP solutions. The major ASSP vendors have been introducing more advanced switch fabric chipsets, with integrated MACs, layer 3-7 functionality (for routing support) and built-in support for uplink ports. Next generation designs incorporate network search engines for more robust routing.

Wireless LAN

Wireless LANs use the **802.11** protocol to send data traffic over unlicensed **2.4 GHz** and **5 GHz** radio waves, mimicking Ethernet but with a wireless connection.

Wireless LAN has seen runaway success in home networking, specifically for distributing a broadband Internet connection wirelessly around the home.

Wireless LAN in a Home Application

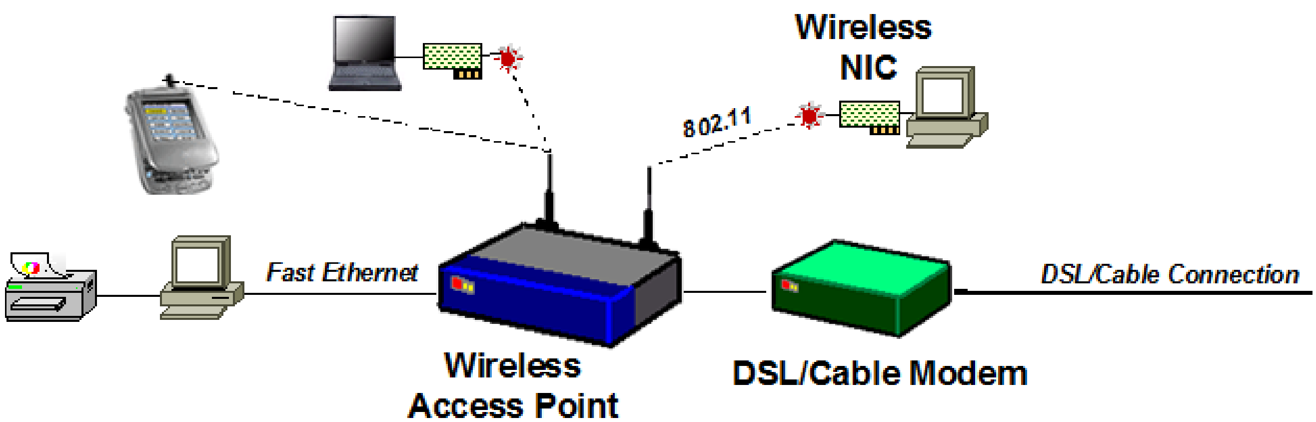

Source: Oppenheimer & Co.

Wireless LANs connect multiple PCs (or other connected devices) in a single location and provide the ability to share files, hardware and resources in a fashion similar to traditional wireline Ethernet LANs. The distinction, of course, is that no physical wiring is needed to connect any single client PC to its neighbor (or to servers, or printers, or PDAs etc.). Instead, wireless LAN encodes data using the 802.11 protocol and broadcasts it over unlicensed spectrum.

Most networking technologies to date have been geared to corporate networks. Wireless LAN, on the other hand, has really been driven by the home networking market. In fact, it is the connected home that has really driven wireless LAN sales to date. In the home setting, 802.11 acts as wireless medium for sharing a broadband Internet connection. Wireless routers designed for the home interface directly with a cable or DSL modem (using Ethernet) and distribute the bandwidth to wireless-enabled PCs.

Wireless LAN is transforming into a home media distribution system, used to send video, voice and data around the home wirelessly. Under this topology, modems, PCs, set-top boxes, game consoles, and VoIP-enabled phones are networked together wirelessly, with the ability to share content and an Internet connection between them. Voice service is distributed wirelessly to WLAN-enabled VoIP phones. Video is distributed from connected PCs, handsets, and digital TVs. Gaming is live with multiple connected participants. The home, and all the devices in it, is now truly connected.

Wireless LAN

In the corporate environment, 802.11 can be used to extend the LAN without wires, allowing for greater mobility among users. We have seen slow but steady uptake as security, bandwidth, handoff, and channelization concerns have been addressed.

Wireless LAN in a Corporate Application

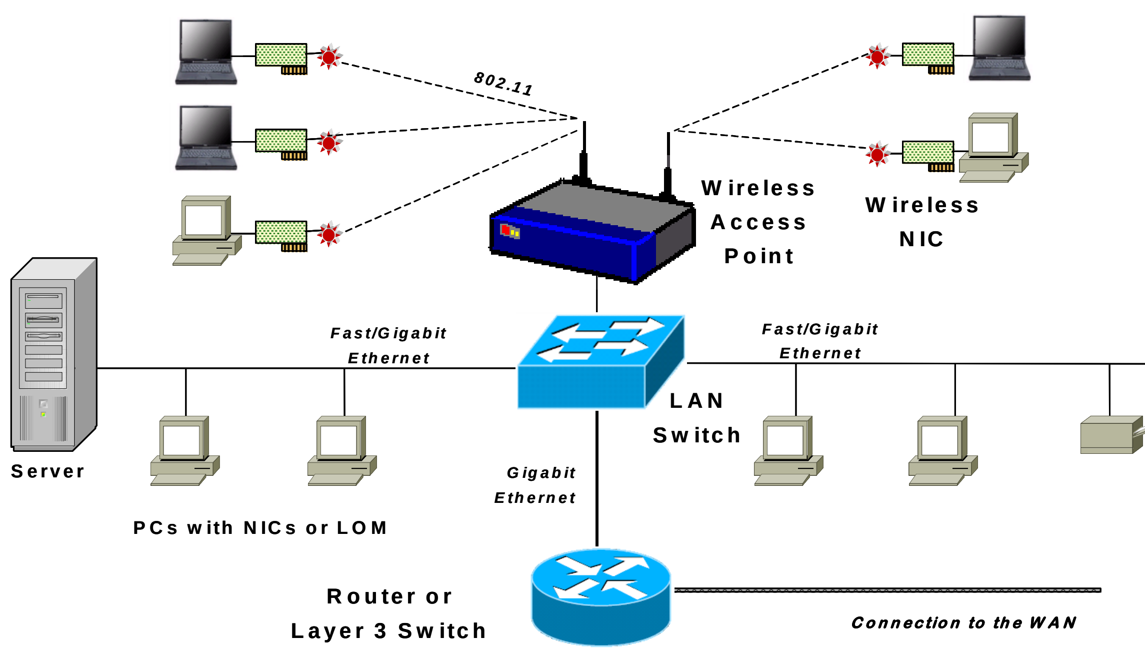

Source: Cisco, Oppenheimer & Co.

In the corporate setting, a WLAN is often meant to augment the connection from user to user currently accomplished with a wired Ethernet LAN. It is generally not designed to replace the traditional LAN switch and backbone architecture, as it requires this equipment to switch and route data and to provide the interface to access and edge routers and out to the WAN. Rather, 802.11 "extends" the LAN without wires. In a typical corporate setting, wireless-enabled PCs access specialized wireless access points distributed at regular geographic intervals throughout the building or campus. Access points act both as transmit/receive bases and as a connection between the wireless and wired LAN.

Wireless LAN technology presents even greater utility when used in a campus setting. Wireless LANs provide true mobility throughout the campus and can be used to provide intranet and Internet access to notebook computers, PDAs and other handheld devices across the campus with no need to connect or configure equipment. They are especially popular with campuses that have inherently mobile users, such as universities and airports.

In terms of adoption, enterprise wireless LAN has been slow to gain steam. This was originally due to IT concerns over issues like security, bandwidth, handoff, and channelization. The major enterprise networking OEMs, however, have made a lot of progress addressing these issues. Tougher security standards have been implemented directly into the 802.11 standard, and computing and networking software have additional safeguards as well. 802.11ac offers more bandwidth and channelization with gigabit speeds. Switch management software and hardware solutions are now coming to market to support dynamic roaming between access points. Power-over-Ethernet is enabling easier access point deployments. We are seeing steady progress in corporate WLAN.

Wireless LAN

Wireless LAN Protocols

802.11b – The first wireless LAN standard to take off, 802.11b runs at 11 Mbps and operates in the 2.4 GHz band.

802.11a – 802.11a runs at 54 Mbps and operates in the 5 GHz band, and offers more channels than 802.11b. It was never a success as a standalone protocol, but is still included in multi-band (a/b, a/g) solutions.

802.11g – 802.11g adds additional modulation schemes to 802.11b, increasing speeds to 54 Mbps while remaining in the 2.4 GHz band. Backward compatible with 802.11b.

802.11n – 802.11n boosts throughput to 108 Mbps or higher. 802.11n uses MIMO technology in both the 2.4 and 5 GHz bands.

802.11ac – Developed from 2011-2013, 802.11ac boosts throughput to 1 Gbps+ and a signal link throughput of at least 500 Mbps. 802.11ac has up to 160 MHz bandwidth and up to eight MIMO spatial streams.

802.11ax – Due for public release in late 2018 – 2019, 802.11ax uses MIMO technology in existing 2.4 and 5GHz bands and expected to have top speeds of ~10Gbps. The protocol improves spectral efficiency, particularly in dense deployments areas

---

Although most wireless LAN technologies use the same topology, they can vary widely in terms of standard, each with its own frequency band, data transfer rate and modulation scheme. A major hurdle holding up volume wireless LAN deployments was the lack of a clear standard for the equipment makers to center around. Until the mid-2000s, two key standards, IEEE 802.11 and HomeRF, were competing for OEM adoption. However, aggressive actions by OEMs and chip manufacturers, together with the formation of the Wireless Ethernet Compatibility Alliance (WECA), which is now know as the Wi-Fi Alliance, helped the industry settle on 802.11 as the dominant standard.

Within 802.11, the first sub-standard to take off was 802.11b, an 11 Mbps technology that operates in the 2.4 GHz band (the same band used by other consumer devices such as cordless telephones, microwave ovens, and Bluetooth devices). This band is free from government regulation, allowing users to set up wireless LANs without license or approval. The FCC has defined three spectrum ranges, known as ISM bands, for use by unlicensed continuous operation up to 1 Watt: industrial (902-928 MHz), scientific (2.4-2.4835 GHz) and medical (5.15-5.35 GHz and 5.725-5.825 GHz). 802.11b was the first wireless LAN standard, operating in the 2.4 GHz ISM band, though it was quickly replaced with the faster 802.11g. 802.11g uses the same frequency—and thus remains backward compatible with 802.11b products—but supports data rates of up to 54 Mbps.

Another sub-standard, 802.11a, offers data transfer rates up to 54 Mbps using the less crowded (but still unlicensed) 5 GHz band. 802.11a received a lot of attention but never much deployment as a standalone technology. The primary advantage of 802.11a, which offers similar data transfer rates as 802.11g, is its higher channel density that provides better average user throughput in congested networks. However, because of its 5 GHz frequency, 802.11a suffers from narrower range and greater power consumption than 802.11g. It is also not compatible with the large installed base of 802.11b and 802.11g. As a result, 802.11a deployment has been included as a second band in multi-band implementations (a/b, a/g solutions), but as mentioned above, very little 802.11a-only equipment ever shipped.

802.11n uses both the 2.4 GHz and 5 GHz bands and leverages MIMO (multiple input multiple output) radio technology to boost bandwidth to 108 Mbps or higher. 802.11n is an effective solution for both the home and enterprise. The first products, which were sold as "pre-n" or "draft-n" solutions, came to market in 2006. Several drafts were approved.

The newest available iteration of the technology is 802.11ac. The standard was developed from 2011 through 2013 with final publication scheduled for early 2014, though 802.11ac is in volume production today in both the home (routers, set top boxes) and handsets. 802.11ac operates on the 5 GHz band and has up to 8 MIMO spatial streams to allow for throughput of 1 Gbps and a single link throughput of at least 500 Mbps. Beamforming with standardized sounding and feedback allows for compatibility between vendors.

802.11ax is due for public release in late 2018 or early 2019, offering up to 10Gbps at top speeds. The protocol uses existing 2.4 and 5GHz spectrum and utilizes MIMO to improve efficiency, particularly in dense, crowded network deployments.

The exhibit below displays WLAN standards and working groups.

IEEE Standards	
802.11a	The first 5 GHz wireless standard, 802.11a operates at speeds up to 54 Mbps. The standard employs a more advanced modulation scheme than 802.11b, called orthogonal frequency division multiplex (OFDM), than 802.11b that allows data to be encoded on multiple radio frequencies concurrently. Products based on 802.11a currently deliver a much shorter range than 802.11b/g because of the higher frequency band, though there is less interference as this band is less crowded.
802.11b	The first standard to gain wideband adoption, 802.11b operates in the 2.4 GHz band at speeds up to 11 Mbps. The standard uses one of several Direct Sequence Spread Spectrum (DSSS) modulation schemes, including Complementary Code Keying (CCK). 802.11b products typically have a range of 200-300 feet. 802.11b has been mostly replaced by 802.11g.
802.11e	The 802.11e standard adds a MAC enhancement for quality of service (QoS) features and multimedia support that are critical for voice, video and audio delivery over wireless LAN. This standard is seen as a critical enhancement toward wide-scale usage in the networked home.
802.11g	802.11g operates at 2.4 GHz, but uses the OFDM modulation scheme for speeds above 11 Mbps to boost data rates. 802.11g-based equipment runs at data rates of up to 54 Mbps, making it competitive with 802.11a. A key advantage of 802.11g is that it is backward compatible with 802.11b; thus, 802.11b NICs will operate with 802.11g access points, though only at 11 Mbps and vice versa.
802.11h	Enhancements incorporated in 802.11h automate frequency band conformance and tuning, employing dynamix frequency selection (DFS) and transmit power control (TPC). The technology enables 5 GHz compatibility in Europe and is also included in HiperLAN/2.
802.11i	802.11i adds security features to improve on WEP and supports 802.1x authentication, Temporal Key Integrity Protocol (TKIP) and advanced encryption standard (AES).
802.11j	802.11j includes extensions for Japan.
802.11n	802.11 supports speeds of 108 Mbps or higher and theoretically can evolve to 540 Mbps over time. 802.11n builds upon previous 802.11 standards by adding MIMO (multiple-input multiple-output). MIMO uses multiple transmitter and receiver antennas to allow for increased data throughput and range.
802.11p	A working group that is extending the 802.11 standard to automobiles. As part of a dedicated short range communications (DSRC) system, 802.11p operates in the 5.9 GHz band and deals with roaming from cell to cell in a fast-moving vehicle.
802.11ac	The newest available iteration of the technology is 802.11ac. The standard was developed from 2011-13 with final publication scheduled for early 2014, though 802.11ac is in volume production today in both the home (routers, set top boxes) and handsets. 802.11ac operates on the 5 GHz band and has up to 8 MIMO spatial streams to allow for throughput of 1 Gbps and a single link throughput of at least 500 Mbps.
802.11ax	802.11ax is due for public release in late 2018 or early 2019, offering up to 10Gbps at top speeds. The protocol uses existing 2.4 and 5GHz spectrum and utilizes MIMO to improve efficiency, particularly in dense, crowded network deployments

Source: IEEE, Wi-Fi Alliance, Wikipedia, TechWeb Encyclopedia, company reports and Oppenheimer & Co.

Wireless LAN

WLAN Equipment Revenue Forecast

- We forecast the WLAN network infrastructure market will grow at a 5% CAGR from $9.5B in 2016 to $11.9B in 2021.
- Enterprise network infrastructure will grow at a 6% CAGR over this period from 57% of total revenue to 61% in 2021.
- Consumer networking WLAN equipment is likely to grow at a more modest 3% CAGR, through 2021. We would expect the majority of consumer WLAN growth to come in wireless devices (smartphones) rather than networking equipment.

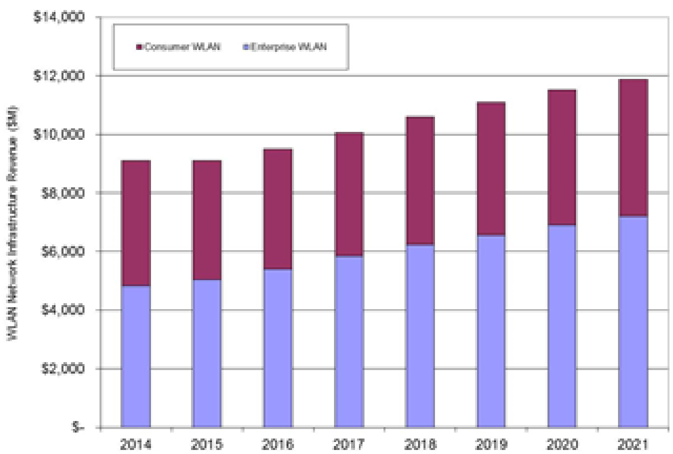

Source: IDC, Oppenheimer & Co.

WLAN Equipment Suppliers

Enterprise WLAN Equipment

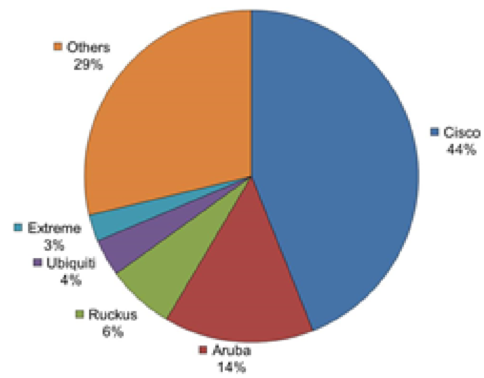

Consumer-Class WLAN Equipment

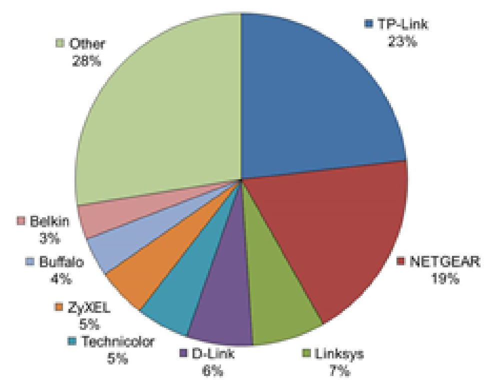

Source: IDC, Oppenheimer & Co.

Wireless LAN Infrastructure Forecast: We forecast continued growth of WLAN infrastructure equipment over the next five years, led primarily by expansion in the enterprise market. We forecast that enterprise WLAN equipment will grow at a 5% annual rate, from $9.5 billion in 2016 to $11.9 billion in 2021. Bring-your-own-device is demanding continued WLAN rollouts in the enterprise infrastructure market and the increasing demand for cellular WiFi offload services by cellular operators is driving further need for enterprise access points. Consumer equipment likely remains growth at a more modest 3% CAGR over the same period, as home networks gain more momentum. Accordingly, consumer access points and gateways in the home continue to reach saturation points in most emerging market and are increasingly tied to slow and maturing upgrade refresh cycles.

Wireless LAN Equipment Suppliers: The top supplier of dedicated enterprise wireless LAN infrastructure equipment (excluding NICs) in 2016 was Cisco with 44% market share. Aruba held the No. 2 spot with 14% share. Ruckus, Ubiquiti, and Extreme rounded out the top five with 6%, 4% and 3% share, respectively. The remaining 29% of the market is highly fragmented with no other supplier commanding more than 2% share of the market.

In the consumer class infrastructure market, TP-Link remained at the top spot with 23%, up from 20% in 2012. Netgear came in second with 19% share in 2016, while Linksys was third with 7% and D-Link held 6% share. Technicolor and ZyXEL round out the top-5 suppliers with each with 5% share.

Wireless LAN

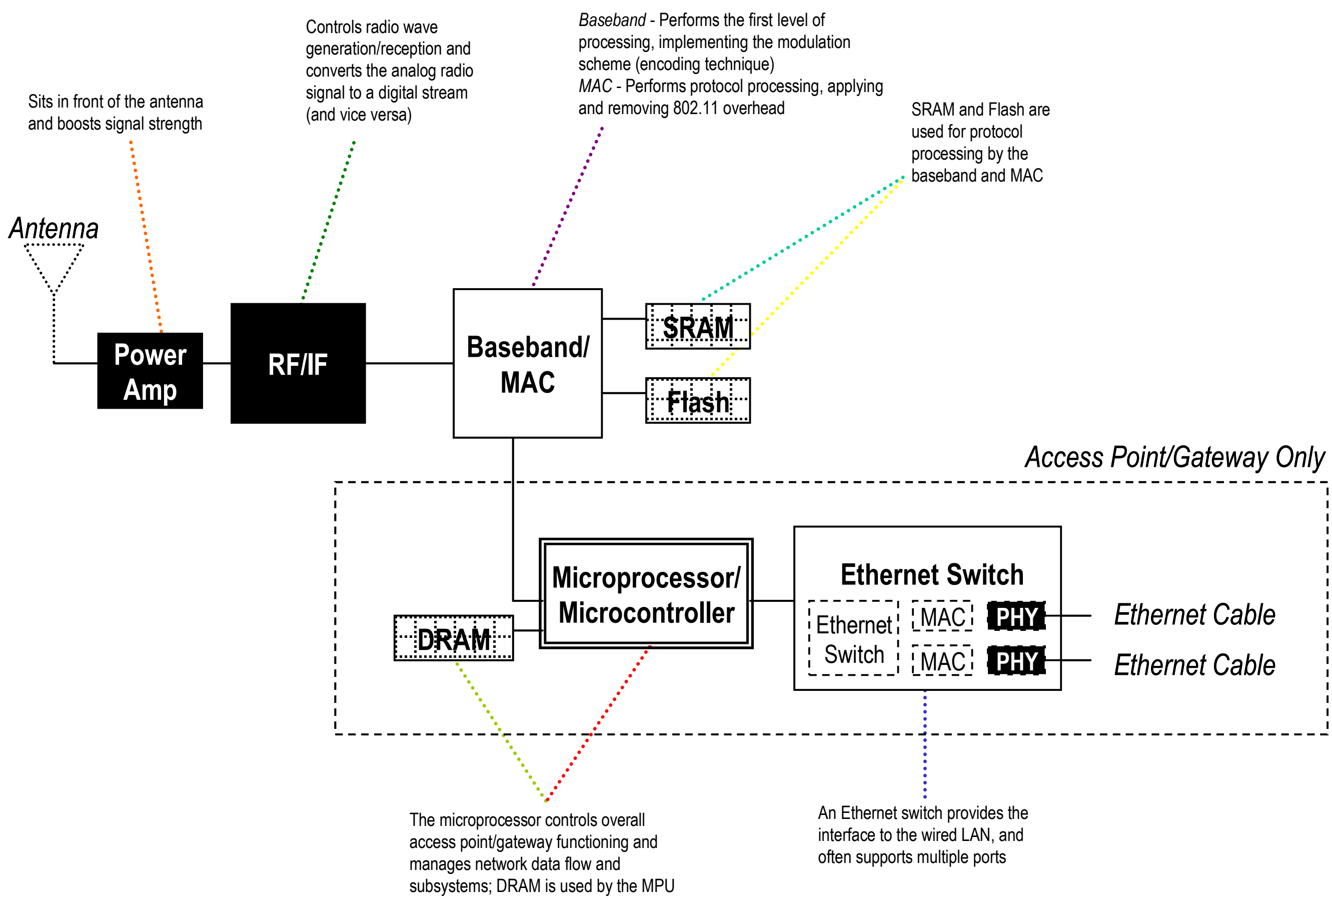

Source: IDC, Oppenheimer & Co.

Wireless LAN client implementations and access points blend devices from the Ethernet and wireless handset worlds. The WLAN signal is received and transmitted through the RF block, which includes a power amp, RF and IF transceiver. The signal is processed by a baseband, which implements the modulation scheme, and then an 802.11 MAC removes or applies 802.11 overhead (the baseband and MAC are usually integrated into a single digital IC). Data packets are then relayed to the host system (in a NIC/client) or to a processor (in an access point or gateway).

Note that the diagram above (and discussion below) shows the core 802.11 chipset (PA, RF/IF, and baseband/MAC), as well as the additional silicon that would be required to enable an access point or a gateway. In a broadband gateway, there would be even more content to enable the modem. Note that our forecast and market share discussion of wireless LAN semiconductors focuses only on wireless LAN-specific ICs, i.e., the 802.11 chipset.

***RF** (Client and Access Point/Gateway)*

- **Power Amp**: The power amp sits in front of the antenna and boosts signal strength.
- **RF/IF Converter (RF)**: The RF front end controls the radio wave generation on the transmit side and receives the radio wave signal on the receive side, translating between low-frequency analog signals used by the system (specifically by the IF transceiver) to high-frequency RF signals used in radio communications. In newer implementations, the RF converter and IF are integrated in a single direct conversion transceiver, as shown in the block diagram on the preceding page.

- **I/Q Modem (IF)**: The IF performs all quadrature modulation of "I" and "Q" baseband signals, converting the analog radio signal to a digital stream to be processed by the baseband and vice versa. In most implementations, the IF is integrated with the RF in a single direct conversion transceiver, as shown in the block diagram above.

Baseband/MAC *(Client and Access Point/Gateway)*

- **Baseband Processor**: The baseband processor performs the first level of digital processing on the signal, implementing the modulation scheme (encoding technique). The baseband and MAC are usually integrated in a digital protocol IC, as shown in the block diagram.
- **MAC**: As in wired Ethernet, the MAC applies and removes overhead and implements the 802.11 protocol. Note that some wireless LAN chipsets implement additional security, QoS, and multimedia functions on top of the base 802.11 standard; the logic that facilitates these features is integrated into the MAC. The MAC is usually integrated with the baseband in a single ASSP, as shown in the block diagram.
- **SRAM and Flash**: These memories are used by the 802.11 MAC.

Gateway Block *(Access Point/Gateway Only)*

- **Processor:** An embedded microprocessor performs control functions in the access point along with some limited data path functions. In most residential gateways they are integrated into the baseband/MAC. Enterprise-class access points tend to have discrete microprocessors, needed to perform advanced access control, security, and management functions not found in consumer wireless LAN.
- **DRAM**: DRAM is the primary memory used by the embedded processor.
- **Ethernet Switch:** Though the basic wireless access point needs only one Ethernet port to interface with the PC or wired network, most access points are actually gateway devices which usually include multiple Ethernet ports. These are usually implemented with a single-chip Ethernet switch, which includes Ethernet PHY, MAC, and switch functions. Most gateways have 4-5 ports of 10/100 Ethernet, though this can vary.

Modem Block *(Broadband Gateway Only – ICs not shown in block diagram)*

- **DSL Line Driver** (used in DSL gateways): A DSL line driver interfaces with RJ-11 telephone line.
- **DSL Data Pump/SoC** (used in DSL gateways): A DSL modem SoC implements the DSL protocol; it has an integrated DSP data pump to implement the DSL protocol and an analog front end to interface with a line driver.
- **Cable Tuner** (used in cable modem gateways): A tuner interfaces with the coax and tunes the cable RF frequency for transmission and reception.
- **DSL Data Pump/SoC** (used in cable modem gateways): A cable modem SoC implements modem functions, and has an integrated front end processing (QPSK/QAM modulation for transmission and QAM demodulation for reception) and a DOCSIS MAC.
- **VoIP DSP:** A VoIP packet processor DSP performs the TDM-to-IP conversion on IP-based phone calls. Integrated codecs perform the analog-to-digital and digital-to-analog conversion.
- **Line Driver/SLIC:** This SLIC is a line driver that interfaces with the telephone line, sending and receiving the analog signals to and from analog telephones.

Wireless LAN

Home WLAN Units

- The WLAN market is increasingly difficult to decipher because of combo chip integration. Standalone WLAN units are in rapid terminal decline in PCs, handsets, tablets and most other consumer devices. Our combo chip forecast is in the wireless section of this report.
- In home routers, gateways and access points, we can witness the evolution of 802.11 technologies over time. 802.11a/b/g/n are the dominant standards today, quickly giving rise to dual band WLAN, which in this case includes 802.11ac.

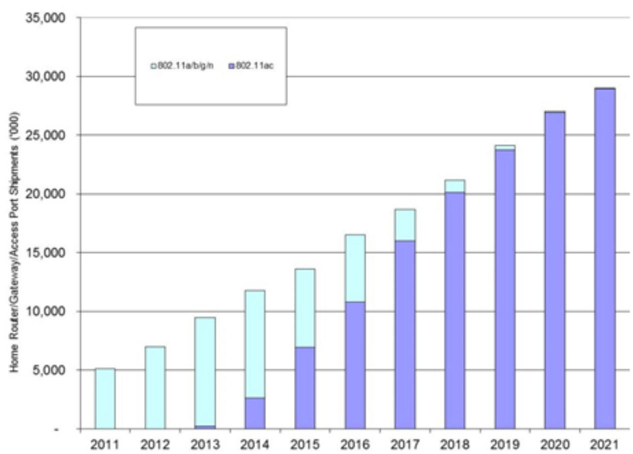

Source: IDC, Oppenheimer & Co.

WLAN Competitors

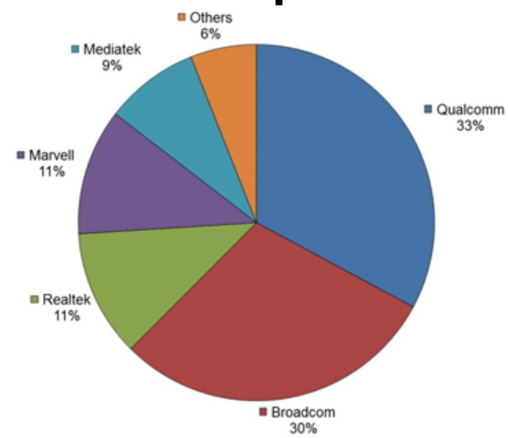

Combo Chip Competitors

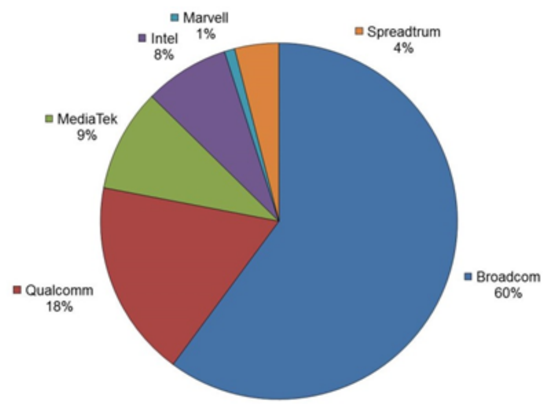

Source: IDC, Oppenheimer & Co.

Wireless LAN Semiconductor Forecast: Wireless LAN is increasingly being offered in combo chips (ex. BT + WLAN + FM) and we believe the total standalone WLAN market is in terminal decline. This affect on the market is twofold; creating a substantially larger unit TAM for WLAN, particularly in smartphones and tablets, but increasingly driving down costs. Aggressive price declines should continue to partially offset unit growth. It is difficult to forecast WLAN units – we forecast mobile smartphone/tablet WLAN units and the total combo chip market in more detail in the wireless section of this report. By technology, we forecast 802.11a/b/g/n to decline at a 61% CAGR through 2021 as 802.11ac and dual-band WiFi begin to aggressively take its place. We forecast dual-band WLAN, which includes 802.11ac, to grow at a 22% CAGR through 2021.

Wireless LAN Semiconductor Competitors: The rapid growth we've seen in wireless LAN has been made possible by the large number of vendors offering competitive solutions. The WLAN market is consolidating around the two largest competitors, Broadcom and Qualcomm. The largest supplier of wireless LAN chipsets in 2016 was Qualcomm, with 33% share (QCOM held 32% in 2014). Broadcom was second with 30% revenue share, down roughly one percentage point from 2014. Marvell, primarily due to exposure in gaming consoles Sony's Playstation platform and Microsoft's Xbox family where Marvell is the sole supplier, and Realtek, a Taiwanese vendor, where tied for third with 11% share. Rounding out the top five was Mediatek with 9% share.

The market share dynamic is noticeably different in combo chip, where Broadcom held 60% market share in 2016 (slightly higher than 57% in 2014). Qualcomm at 18% has aggressively ramped its platform following Texas Instruments exit from the business several years ago. Chinese vendor MediaTek held 9% share in 2016, after not having a presence in 2012.

Wireless LAN

Wireless LAN Semiconductor Market Trends

•The market began to consolidate in 2011; Qualcomm purchased Atheros and Mediatek acquired Ralink

•802.11ac has quickly become the dominant technology, accounting for 65% of enterprise port shipments in 2016. 802.11ac port shipments are expected to grow at ~30% CAGR the next several years while legacy platforms decline rapidly. 802.11ac is expected to account for nearly 100% of the market in 2-3 years.

•WLAN has reached maturity in the PC market.

•For chipset vendors, the handset and tablet markets are clearly where the growth opportunities are considering unit volumes, pricing and replacement rates

•Integration into combo chips is becoming standard. Full wireless SoC integration is next.

- **The market began to consolidate in 2011; Qualcomm purchased Atheros and Mediatek acquired Ralink. Both are share gainers.** As connectivity becomes an increasingly large theme in the global economy, large semiconductor vendors quickly realized the benefits of bringing advanced WLAN capabilities in house. Qualcomm acquired Atheros, MediaTek acquired Ralink and both vendors gained market share in both WLAN and connectivity combo chips. That traditional baseband vendors are positioning for the growth in WLAN and combo chips points to the importance of connectivity going forward, in addition to the evolution into a wireless SoC over time.

- **802.11ac has already become the dominant wireless protocol.** 802.11ac has quickly become the dominant technology, accounting for 65% of enterprise port shipments in 2016. 802.11ac port shipments are expected to grow at ~30% CAGR the next several years while legacy platforms decline rapidly. 802.11ac is expected to account for nearly 100% of the market in 2-3 years.

- **WLAN has reached maturity in the PC and smartphone markets.** Wireless LAN is a standard feature in notebooks. While desktop penetration is likely to continue to creep higher through 2015/2016, the WLAN PC market as a whole has reached the point of maturity, growing (or declining) roughly in-line with total notebook units, and growing roughly in-line with handset and tablet units.

Integration into combo chips is becoming standard. Full wireless SoC integration is next. Chip vendors have had astounding success in the smartphone market with combo chip integrating Wi-Fi on the same chip as Bluetooth, GPS and FM. Not only does this lower cost, but also size and allows smartphone vendors to move quickly to market with new designs by sourcing all connectivity demands from a single vendor. QCOM has also been active in the integration of RF functions, incorporating GPS and Bluetooth with its basebands. WLAN has become close to entirely-combo chip already. The next leg of integration will likely come in the form of a full wireless SoC; baseband + application processor + combo connectivity. QCOM has already integrated a WiFi control chip into its Snapdragon processors and virtually the entire market is moving to an integrated solution.

Bluetooth

Bluetooth is a **personal area networking** technology, used to send data between devices in close proximity at rates of up to 50 Mbps.

Bluetooth has found its primary volume application in the handset market, where it functions as a reliable protocol to connect the handset to other devices, including wireless headsets, PCs, and consumer devices. PC peripherals and consumer portables are also using Bluetooth as a cable replacement for connection to the PC, or to connect to one another without the PC.

Bluetooth is one of the leading personal area networking technologies, a short-range, low-power, high-bandwidth wireless transmission protocol aimed at connecting wireless handsets, PC peripherals and consumer devices. Bluetooth is standards-based, universal and cost effective, offering relatively high bandwidth (up to 50 Mbps in the latest Bluetooth 5 standard) over a short distance (usually 30-300 feet, but up to 800 depending on the device class) in an unregulated frequency band (2.4 GHz), with extremely low power consumption and low implementation cost.

Bluetooth has found its primary volume application in the wireless market, where it functions as a reliable protocol to connect the handset to other devices, including wireless headsets, automotive stereo systems, PCs and consumer devices. Smart handheld devices, PC peripherals, and consumer portables are also using Bluetooth to perform data transfers without wires; and in many cases the entire PC can be removed from the equation (e.g., a digital camera can print straight to a Bluetooth printer, a smartphone can transfer video directly to a connected TV, etc.). Game consoles also use Bluetooth to communicate between console and controller.

In terms of protocol, Bluetooth version 1.2 supported a data transfer rate of 1 Mbps, but was replaced by Bluetooth 2.0 and now, version 3.0. The Bluetooth 3.0 link is used for negotiation and establishment of the signal, but the actual data traffic is carried over a co-located 802.11 transfer. Bluetooth 4.0 includes Classic Bluetooth, Bluetooth High Speed and Bluetooth Low Energy, but is commonly referred to as BLE, or Bluetooth Low Energy. Bluetooth 5 offers up to 50Mbps and up to 800 feet range, recently introduced in the flagship Samsung Galaxy S8 and IPhone 8/8+/X.

Bluetooth is rapidly being integrated into combo chips in most devices. We therefore will not provide a Bluetooth market forecast or market share of competitors. Please refer to the wireless section of the report for our combo connectivity forecast.

Bluetooth

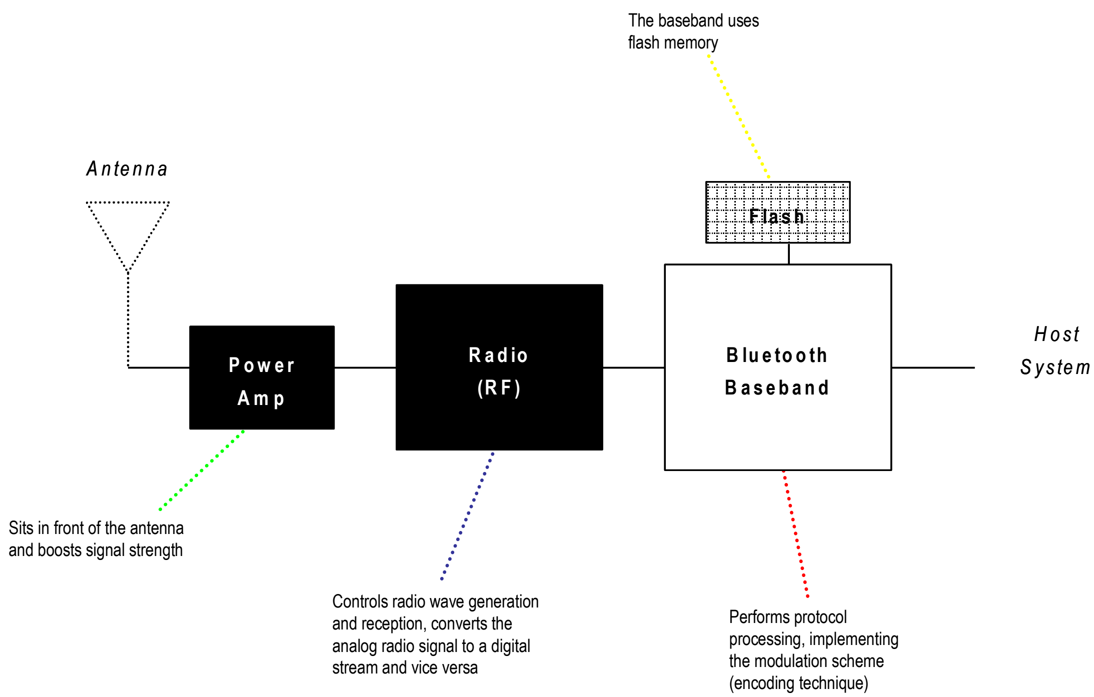

Source: Oppenheimer & Co.

Bluetooth modules have similar configurations to those of wireless LAN client implementations. A power amp interfaces with the antenna and relays signals to a radio (RF) chip. The radio chip processes the RF signal and passes it to the Bluetooth baseband, which implements the Bluetooth modulation scheme (the encoding technique) and extracts the data packet. On the transmit side, the process is reversed: the baseband encapsulates the data using the Bluetooth modulation scheme and sends it through the RF and power amp to the antenna.

Note that most Bluetooth solutions combine all functions on a single chip, sometimes leaving out the power amp but still integrating the RF and baseband. Other implementations use just a Bluetooth RF, integrating baseband functionality with a host processor (e.g., Qualcomm integrates the Bluetooth baseband in its handset baseband IC).

- **Power Amp:** The power amp sits in front of the antenna and boosts signal strength.
- **Radio (RF)**: The RF controls radio wave generation on the transmit side and receives the radio wave on the receive side.
- **Baseband Processor**: The baseband processor implements the Bluetooth modulation scheme, extracting the packet from the Bluetooth frame on the receive side and encapsulating in on the transmit side.
- **Flash**: The baseband usually uses a small quantity of flash memory.

Storage

The buildout of high-performance computing infrastructure, adoption of higher-bandwidth LAN and WAN connections, new applications offering richer content, and increased focus on data security and redundancy have driven continued strong demand for high-capacity storage in the enterprise. Petabyte shipments continue to grow at a double-digit (%) CAGR Y/Y.

Storage systems are evolving to meet ever-rising needs of IT managers. More complex direct-attached RAID systems and networked topologies like **SAN** and **NAS** offer greater scalability, manageability, and performance for the modern enterprise.

The three preceding networking technologies we have discussed in this section are all used in the transference of data between computing devices. We now turn from data networking to enterprise storage, which is used “behind the scenes” in the enterprise network for the storage, retrieval, backup, and management of corporate software and data.

Over the past few years, storage equipment makers have attempted to bring the benefits of networking to enterprise data storage. The legacy model, which stores applications and data files on disk drives within individual servers, is being replaced by newer networked storage systems. By treating storage as a separate resource from processing, enterprise IT managers are realizing greater efficiency, reliability, scalability and upgradeability in their infrastructure.

There is a wide variety of storage topologies and technologies employed in enterprise IT systems today. Most systems use a direct attached storage (DAS) architecture, which connects servers directly to drives and drive arrays. Though they do not offer the flexibility, scalability, manageability, and performance of networked storage, DAS systems are robust and easy to set up and maintain, and therefore remain an important part of the IT infrastructure.

High-performance and high-reliability storage systems use storage area networks (SAN), predominantly based on the Fibre Channel protocol. These mimic the LAN topology, using specialized storage switches and adapter cards. SANs are the topology of choice for the highest-end storage applications.

Some simpler IT systems use network-attached storage (NAS), which connect drive appliances directly to the Ethernet network. This brings the benefits of networked storage while leveraging the LAN (no new SAN network is needed).

Storage

Direct Attached Storage

Direct attached storage systems are primarily implemented with RAID controllers, which are directly connected to disk array systems or internal drives. The cheapest, simplest, and easiest to implement, direct attached is the most widely used storage topology.

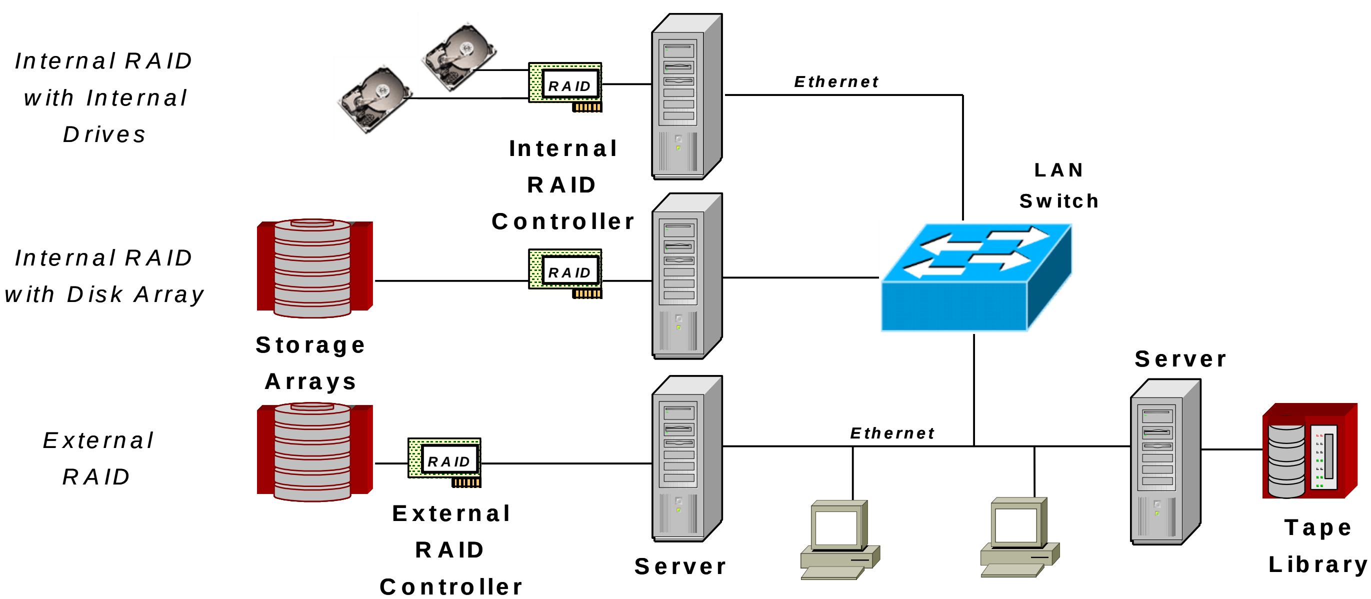

Source: Oppenheimer & Co.

Direct attached storage systems are the cheapest, simplest and easiest to implement and are currently the most widely used storage topology. In a DAS system, servers are connected via standard hard disk drive cables either to enterprise-class hard disk drives within the server or to storage arrays that contain multiple drives. The interaction between the server and the disks is controlled by RAID (redundant array of independent disks), a set of protocols for data striping, duplication, and drive management. A RAID controller, which sits between the server and the individual disks, implements the RAID software stack to enable these features, and allows the server to access an array of hard disk drives as if it were a single, huge drive.

RAID controllers can be either internal or external. In an internal configuration, the RAID controller is housed inside the server and connects to drives within the server or to simple disk arrays like JBODs (just a bunch of disks). In an external configuration, the RAID controller is housed inside the disk array itself, and the server connects to the disk array over a standard hard disk drive cable.

As mentioned above, DAS systems are cheap, simple and easy to implement. The multiple flavors of the RAID standard allow for varying degrees of redundancy and performance enhancement. However, they have limitations in terms of reliability, scalability and manageability that make DAS systems difficult to support as capacity grows larger. SANs have made major headway against large DAS systems over the past few years, especially for the highest-end storage applications.

Storage

Storage Area Networking (SAN)

In a SAN environment, separate Fibre Channel networks switch storage traffic between disk arrays and servers. Storage is managed as a single resource, and tasks such as backup and disk maintenance can often be performed without any processing by host servers or PCs.

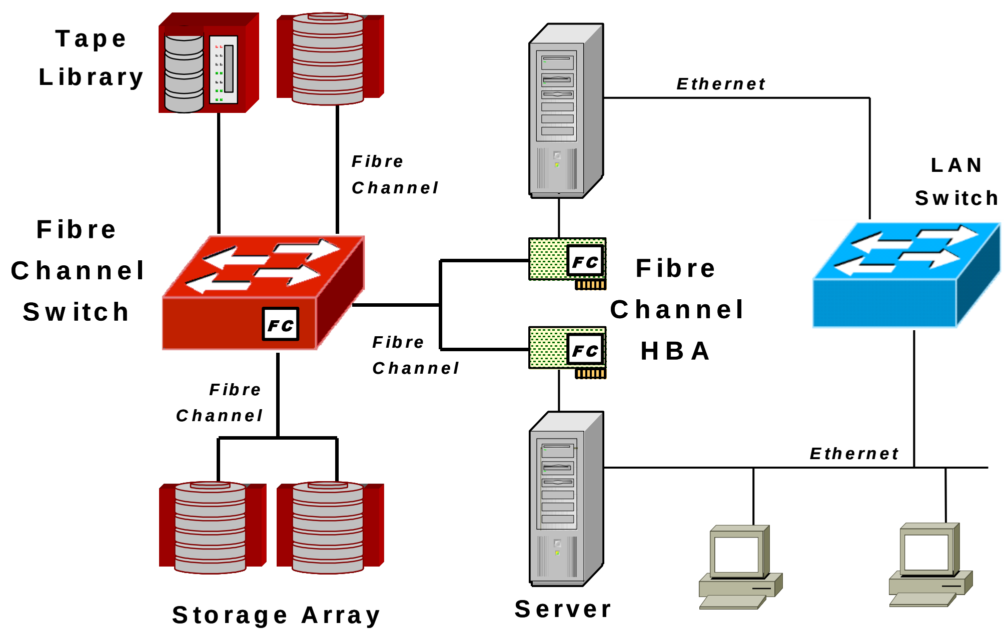

Source: Oppenheimer & Co.

Storage area networks (SANs) are implemented with disk arrays sitting on a high-speed storage-focused network separate from the data network. In this way, processing functions are separated from storage, which is well-suited to large corporate networks. SANs allow IT managers to treat all storage as a single resource and also allow tasks such as backup and disk maintenance to be performed without any processing by host computers.

In a SAN, servers do not connect directly to drives or arrays; they instead connect to a storage-specific network designed to carry storage traffic using the Fibre Channel protocol. These storage networks serialize SCSI traffic into Fibre Channel frames for high speed transport. Fibre Channel is somewhat similar to Ethernet, but is optimized for mass transfers of data as opposed to IP packet routing. The protocol uses optical fiber interconnects and runs at up to 12 Gbps.

To connect to the SAN, each server is outfitted with a Fibre Channel host bus adapter (HBA), which connects to a network of specialized Fibre Channel switches. These switches provide the connection between servers and storage arrays, which also have to be specialized to run over Fibre Channel. Tape libraries can also be connected to the SAN, allowing for backup to take place without the use of servers on the data network.

Note that as SANs have evolved, OEMs have begun building more intelligent appliances to help manage the network such as SAN routers which interconnect between islands of storage in disparate locations. As these are used in only the highest end implementations, we do not picture them in our network diagram.

Storage

Network Attached Storage (NAS)

In a NAS, storage appliances are connected directly to the Ethernet local area data network. Instead of a full-blown operating system, NAS systems use a slim microkernel OS specialized for handling only file reads and writes. NAS systems operate like mini-SANs in that they are self contained storage servers.

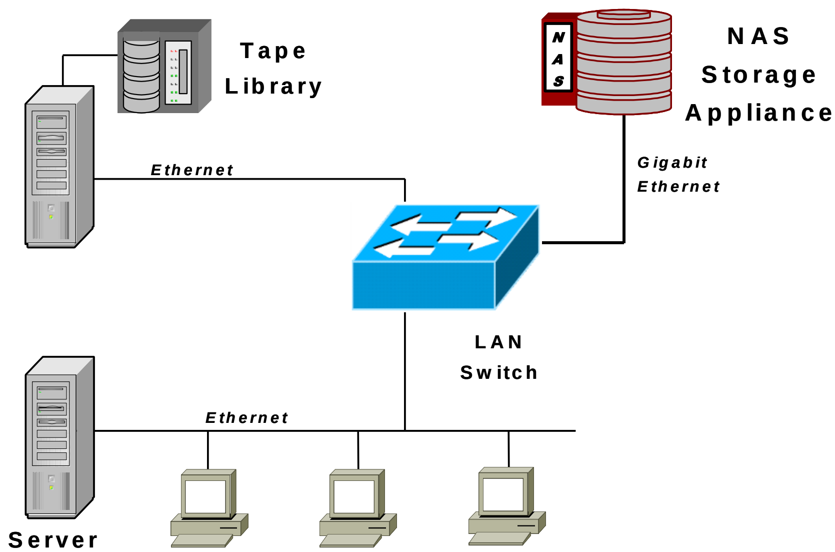

Source: Oppenheimer & Co.

In a NAS environment, storage appliances are connected directly to the Ethernet LAN. NAS appliances operate like mini-SANs in that they are self-contained storage servers. They can be added to the IT infrastructure simply by plugging them into the Ethernet LAN, like any other Ethernet client.

NAS systems usually contain a number of hard disks, often arranged into logical, redundant storage containers or RAID arrays. Instead of a full-blown operating system, NAS systems use a slim microkernel OS specialized for handling only file reads and writes. Storage capacity is added by connecting NAS appliances to the Ethernet network as simply as if a new PC or server were being added.

Note that NAS technology has begun penetrating the consumer market in recent years, sparked by the rise of digital content and home networks. Home NAS appliances can do more than simple USB external HDDs, as they allow for direct access by more than one PC or connected device on the network, and can also connect to other drive arrays for automatic backup. Note that Gigabit Ethernet was an important enabler for this market, as 10/100-based devices were dramatically outpaced by USB 2.0 at 480 Mbps or Firewire 800 in terms of raw data transfer. Today, USB 3.2 boasts speeds of up to 20Gbps; however, the technology is primarily used for connecting to client storage devices such as a NASD or external HDD, rather than enterprise networks.

Storage

Types of Storage Equipment

Enterprise Hard Disk Drive - Enterprise hard disk drives are the highest-performance drives on the market. They are usually housed either inside servers, in storage arrays, or inside NAS appliances, and use the SCSI or SAS protocol. Enterprise drives are included in our forecast for HDDs.

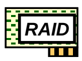

RAID Controller – A client card used to connect servers to storage arrays; they can be housed either in the server (internal) or in the array (external).

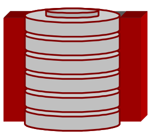

Storage Array - Storage appliance with multiple hard disk drives used in enterprise storage applications. In DAS systems, they either incorporate RAID controllers or connect to internal RAID controllers inside servers. In SAN systems, they connect to Fibre Channel switches.

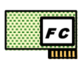

Fibre Channel HBA - A client card used to connect servers to the Fibre Channel network. Analogous to a NIC in an Ethernet network.

Fibre Channel Switch - The basic SAN building block, Fibre Channel switches connect and switch storage traffic between servers and storage arrays.

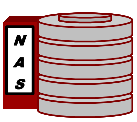

NAS Appliance - NAS appliances integrate storage and LAN connectivity. They include multiple HDDs and a slim OS to manage data transfers.

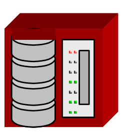

Tape Library – Tape library appliances are optimized for backup functions in an enterprise IT infrastructure system. They are implemented independently of the storage topology.

Types of Storage Equipment

Enterprise Hard Disk Drive: Similar to HDDs used in PCs, enterprise hard disk drives are the highest performance drives on the market. They are usually housed either inside servers, storage arrays, or NAS appliances. Hard disk drives used in storage systems are not included in our forecast, and are instead discussed in the section on HDDs.

RAID Controller: A client card used to connect servers directly to storage arrays. They can be housed either in the server (internal) or in the array (external), and use standard hard disk drive interfaces like SCSI, ATA, SATA, or SAS.

Storage Array: Storage appliance with multiple hard disk drives used in enterprise storage applications. In DAS systems, they either incorporate RAID controllers or connect to internal RAID controllers inside servers. In SAN systems, they include Fibre Channel connectivity to connect up to SAN switches.

Fibre Channel HBA and Switch: An HBA is a client card used to connect servers to the Fibre Channel network. A Fibre Channel switch is the basic SAN building block used to connect and switch storage traffic between servers and storage arrays.

NAS Appliance: A mass storage array optimized for NAS systems, NAS appliances integrate storage with Ethernet LAN connectivity and can be added directly to the Ethernet network to provide enterprise storage.

Tape Library: Storage appliance optimized for backup functions in an enterprise IT infrastructure system.

Storage

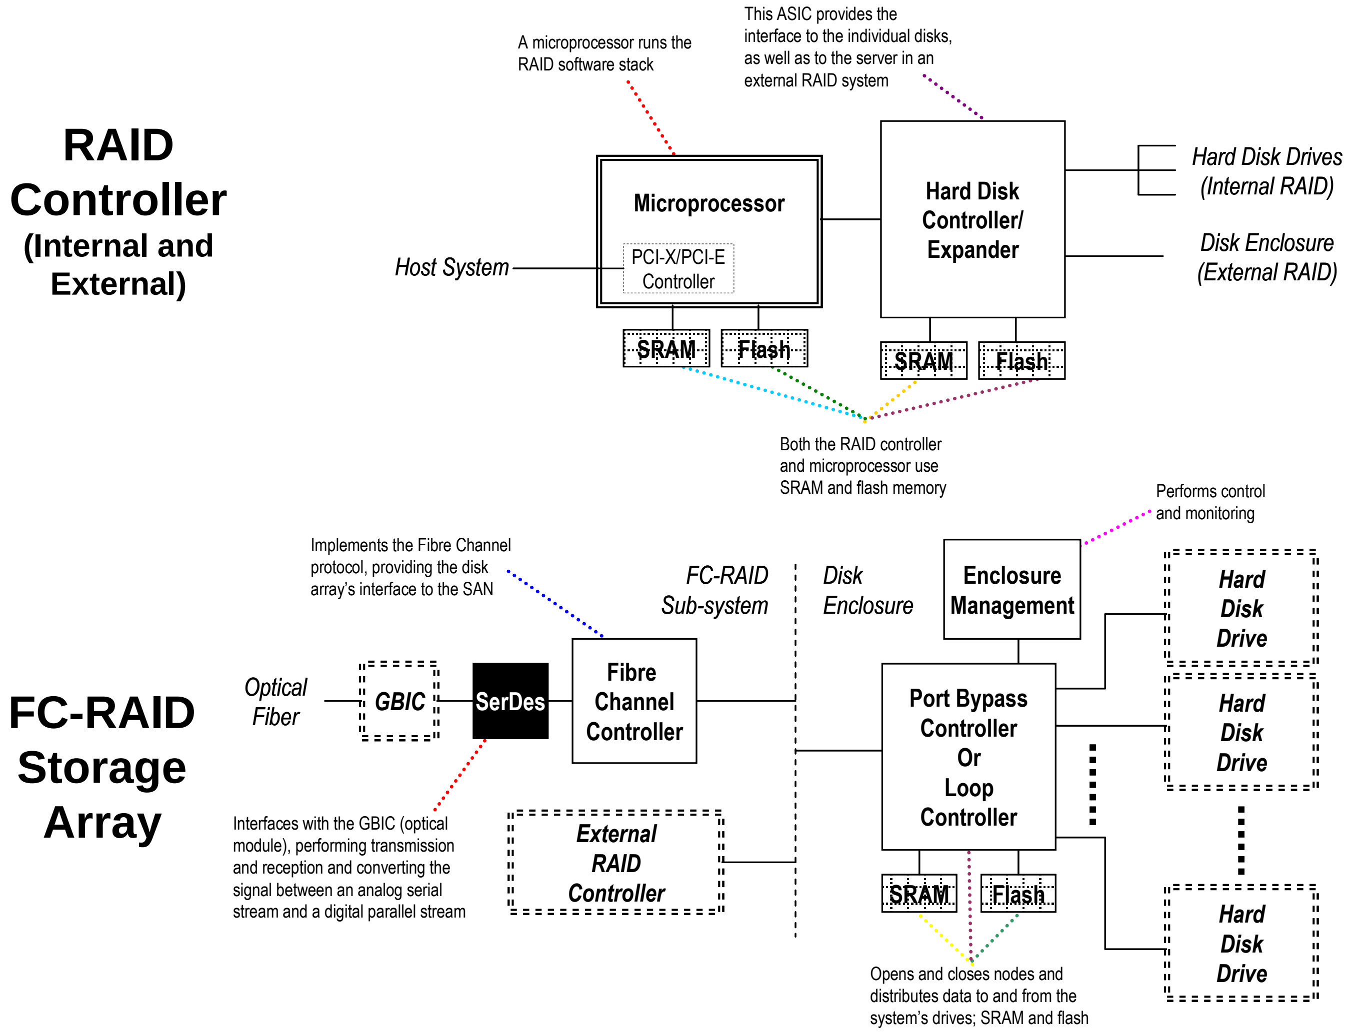

Source: IDC, Oppenheimer & Co.

RAID systems enable critical storage functions such as data redundancy, striping (for performance), and disk management. Their architecture and implementation vary widely, from low-end systems that control just a few disks within a single server, targeted for small/medium business, all the way to large switched disk arrays managing thousands of drives used in large corporate enterprises.

At the core of a RAID system is the RAID controller, which performs the compute-intensive calculations needed to implement the RAID stack at high speed. A RAID controller is usually implemented as a separate add-in board. In internal RAID systems, the RAID controller is used directly inside the server, and is attached directly to the drives, which are also located within the server.

In an external RAID system, the RAID controller is implemented within an array of disks as a subsystem to perform RAID processing. Larger disk arrays will usually have a number of subsystems, including a Fibre Channel subsystem to interface with a Fibre Channel SAN and a disk enclosure to manage the communication with individual disks.

RAID Controller – Internal *(used in servers)*

- **IOP/Microprocessor**: A microprocessor runs the RAID protocol stack, enabling the striping, redundancy, and disk management features of RAID. The microprocessor in an internal RAID implementation is sometimes called an I/O processor (IOP). These processors are usually optimized for high-speed RAID processing, and usually include the interface to the host controller, which is usually PCI-X or PCI-Express based.
- **Hard Disk Controller:** The hard disk controller provides the interface to the individual drives.

- **ROC (RAID-on-Chip**, not shown in above diagram): A ROC chip includes both the IOP and hard disk controller. ROC solutions can be used either on discrete RAID controller cards, but are mainly targeted for ROM (RAID-on-motherboard) implementations.
- **SRAM and Flash**: SRAM and flash are used by the microprocessor as well as by the hard disk controller ASIC.

RAID Controller - External *(used within disk arrays)*

- **Microprocessor**: A microprocessor runs the RAID protocol stack, enabling the striping, redundancy, and disk management features of RAID. In higher end systems, multiple high-performance processors are sometimes used.
- **Hard Disk Controller:** The hard disk controller performs the SCSI/ATA processing to communicate with the drives within the enclosure. In lower end PATA/SATA/SAS RAID systems that do not have separate interface cards, this controller also interfaces with the server on the front end. In higher end systems (especially Fibre Channel), this is not needed since a separate controller card performs this interface.
- **SRAM and Flash**: SRAM and flash are used by the microprocessor as well as by the hard disk controller ASIC.

FC-RAID Storage Array

- **SerDes:** The SerDes interfaces with the optical module, performing transmission and reception and converting the signal between an analog serial stream and a digital parallel stream.
- **Fibre Channel Controller**: A Fibre Channel controller ASIC implements the Fibre Channel protocol, providing the interface to the SAN.
- **Port Bypass Controller (PBC)**: The primary logic device in the disk enclosure, a PBC opens and closes nodes between drives, distributing data to and from the system's drives.
- **Loop Controller**: More complex Fibre Channel disk arrays use loop controllers for a switched architecture instead of a PBC.
- **Enclosure Management**: In high-end systems, enclosure management chips monitor and control the drives in the storage array.
- **SRAM and Flash**: SRAM and flash are used by the port bypass controller or loop controller.
- **Expander** (used in SAS/SATA-based arrays, not shown in the diagram)**:** An expander IC, used in serial arrays like SAS or SATA (instead of a PBC or loop controller), increases the number of drives that can be included in the array. Note that SAS expanders can support both SAS and SATA drives; therefore, most IC designers are focusing on SAS expanders.

Storage

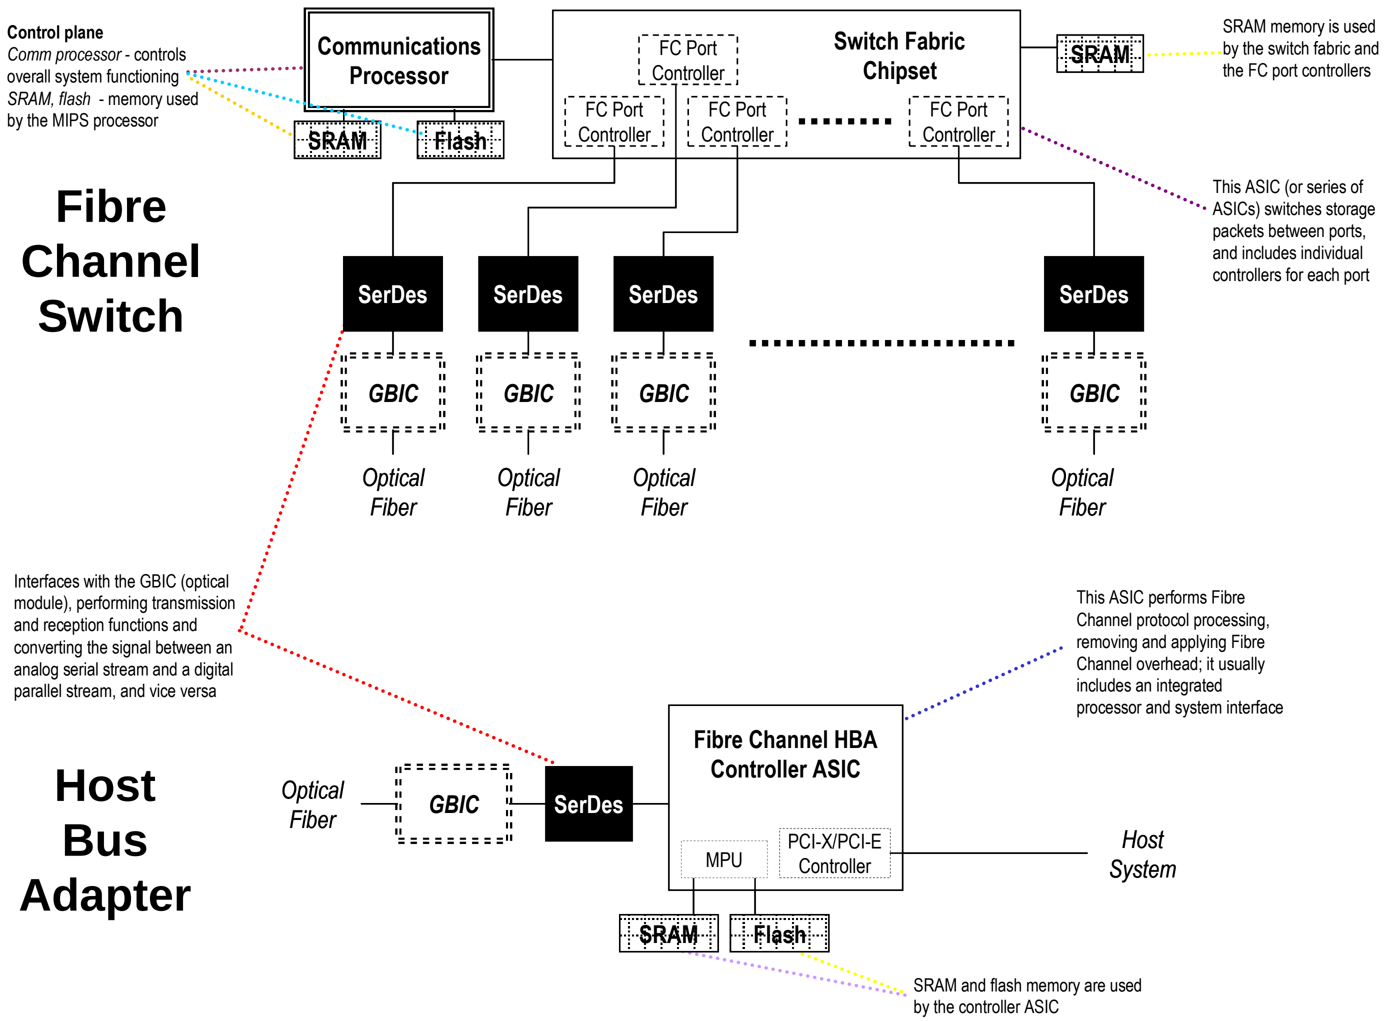

Source: IDC, Oppenheimer & Co.

Fibre Channel Switch

- **SerDes**: The SerDes converts the serial analog stream from the GBIC to a digital parallel stream and vice versa.
- **Switch Fabric Chipset**: The switch fabric chipset is a set of ASICs that switch storage traffic between switch ports. The chipset usually has integrated port controllers that interface with the SerDes on the front of the switch.
- **Communications Processor, SRAM and Flash**: A processor, usually MIPS or PowerPC based, performs control plane functions. The processor will usually use SRAM and flash, while some additional SRAM is also used by the switch fabric chipset in switching and port control.

Host Bus Adapter (HBA)

- **SerDes**: The SerDes converts the serial analog stream from the GBIC to a digital parallel stream and vice versa.
- **Fibre Channel Controller ASIC**: The controller ASIC performs protocol processing on the Fibre Channel stream, applying and removing Fibre Channel overhead to storage packets. The controller usually incorporates a processor, though the HBA can have a discrete MPU as well. It also usually includes a host interface.
- **SRAM and Flash**: SRAM and flash are used by the controller ASIC and the embedded microprocessor.

Telecom/Datacom

We now move from networking equipment used in singular locations to wide area communications technologies used to cross long distances. The WAN blends legacy equipment initially designed for voice telecom with newer technologies designed for data. Although the transmission media used is similar, the technologies and equipment that sit behind it are different: voice networks are generally **circuit-switched** while data networks are generally **packet-switched**.

The wide area network can be thought of in three segments: **long-haul**, used for city-to-city or country-to-country transport, **metro**, used for intra-city transport, and **access**, used to connect individual enterprise and residential customers to the network. In each case, the traditional voice infrastructure has been overlaid with data networking equipment used to switch, route, and transport data alongside voice.

We now move from networking equipment used in singular locations to the telecom and datacom realm, the "wide area networks" (WANs) that move voice and data traffic over longer distances. These networks are a lot older than the data-centric networks we have discussed so far; they are also a lot more challenging to design, build, and maintain. An extra layer of complexity is added with the blending of legacy technology originally designed specifically for voice telecommunications with newer technologies designed to send data traffic over the public telecom network.

Voice Networks

Wide area networks began as strictly voice networks, designed to offer local and long distance telephone service. As such, they emphasize the requirements of high-reliability voice telecommunications. In these networks, quality-of-service is key, as tolerance for latency, information loss and variations in speed of delivery are extremely low in a voice call. Further, information must be delivered strictly in the order it is received and it must be able to flow in both directions simultaneously (known as full-duplex). On the other hand, bandwidth requirements for voice are minimal and usage patterns are stable and predictable. Thus, while voice networks have historically been expensive to install and run, they can be operated at close to full capacity, reducing the need for extensive headroom in the network.

Voice networks have historically been circuit-switched; the transmission is enabled by establishing a dedicated communication channel (a closed circuit) between devices on a network. Such channels are established on demand and as available and provide continuous and exclusive access between the terminal devices until the connection is terminated. In other words, upon initiation of a telephone call, a dedicated channel through the network and switching matrices is established and a prescribed level of bandwidth is guaranteed for the duration of the connection.

Data Networks

With the growth of Internet data technologies such as e-mail, Internet, and e-commerce, service providers have made significant investments to layer data network technologies on top of their voice infrastructure. Though the expertise and expense required to run data over voice-centric networks are significant relative to pure data networks, providers saw this as a better option than building entirely new data-only networks. We therefore find a public data network with a blend of voice and data equipment.

In general, data traffic travels over the same pathways that voice traffic does in the WAN, though the equipment that handles data traffic at each point-of-presence (POP) may be different. Data networks do not require the same quality-of-service that voice networks do, and in fact data traffic does not necessarily need to be delivered in the same order it was sent. On the other hand, bandwidth requirements are large and extremely variable (often referred to as "bursty") and usage patterns are very unpredictable, requiring lots of headroom in the network.

In general, data traffic is better suited to packet-switched networks, which involve encapsulating data in fixed- or variable-length packets and sending them across a shared network. Each packet contains headers encoded with a destination address and is switched at each node in the network over the most appropriate and available physical circuit. No direct circuit or channel is established between sender and recipient as in a circuit-switched network; instead, data packets bounce around in the network until they reach their destination addresses.

That is the big difference between circuit-switched and packet-switched networks. Packets travel through the network independent of each other, and in fact, multiple packets sent from one sender addressed to the same destination may travel different physical paths, depending on the availability of network segments and the level of congestion across each network segment. As the route through the network is constantly changing, different packets may meet with varying levels of delay during transmission through the network. As a result, packet switching does not guarantee a prescribed level of bandwidth and packets are transmitted on a "best efforts" basis only; however, they are generally more efficient, as equipment is shared across a greater number of transmissions and as less "dead air" is transmitted.

Network Segments, Speeds and Protocols

The wide area network is usually thought of in three segments: long-haul, used for city-to-city or country-to-country transport; metro, used for intra-city transport; and access, used to connect enterprise and residential customers to the network. The protocols used to traverse these networks vary somewhat, though there is a lot of overlap. In each case, the traditional voice infrastructure (which of course remains in use for voice) has been overlaid with data networking equipment used to switch, route and transport data such as e-mail, Web, and VPN. In the case of the access network, traditional cable video networks have been overlaid with data equipment in a similar fashion.

In general, the closer to the core (long-haul portion) of the network, the faster the transmission speed. This is because fewer fiber connections are laid between cities than within them and thus the level of aggregation—which translates into higher bandwidth requirements and faster speeds—is greater for networks that span longer distances.

Long-haul networks typically employ transport equipment that can transmit at 10 Gbps (OC-192), 40 Gbps (OC-768), 100 Gbps (OC-3072), or even 200Gbps (OC-3840) and use the SONET/DWDM protocol for transport over optical fiber.

Metro networks also use SONET for transport, usually at speeds up to 10 Gbps (OC-192); however, 100G metro buildouts have also begun. Carriers also use these T/E lines to connect to wireless basestations to reach cellular customers, and for large enterprise customers that want a direct connection to the WAN. Running on top of these SONET and T/E transport networks might be other protocols such as ATM, which is used for multi-service switching of data, voice, and video.

The access network is a bit different, as it casts a much wider net and is used to reach individual voice and data customers. Speeds and protocols include 64 Kbps analog voice (over copper wire), 56 Kbps analog modems (over the same copper wire), 384 Kbps-24 Mbps DSL (again over the same copper wire but using separate frequencies), 1-30 Mbps cable modems (over coaxial cable), or newer 10G PON/FTTH (over fiber).

Depending on the source, destination and access technology, information sent over the WAN may employ a number of protocols in order to reach its destination. For example, an e-mail sent from our Oppenheimer offices to a friend on the West Coast will first exit our corporate LAN via an access router, travel over a T1/T3 line in Verizon's (or another LEC's) network, be aggregated into an OC-3 or OC-12 SONET signal for transport to a larger network node, where it will be further aggregated (multiplexed) into a higher speed OC-48 or OC-192 SONET signal for transport over Sprint's (or some other IXC's) long-distance network. Once it reaches the West Coast, the SONET signal will be disaggregated (demultiplexed) into a lower speed SONET signal in AT&T's (or another LEC's) network, and switched on an ATM network into a DSL signal for transport to the recipient.

Telecom/Datacom

The Voice Network

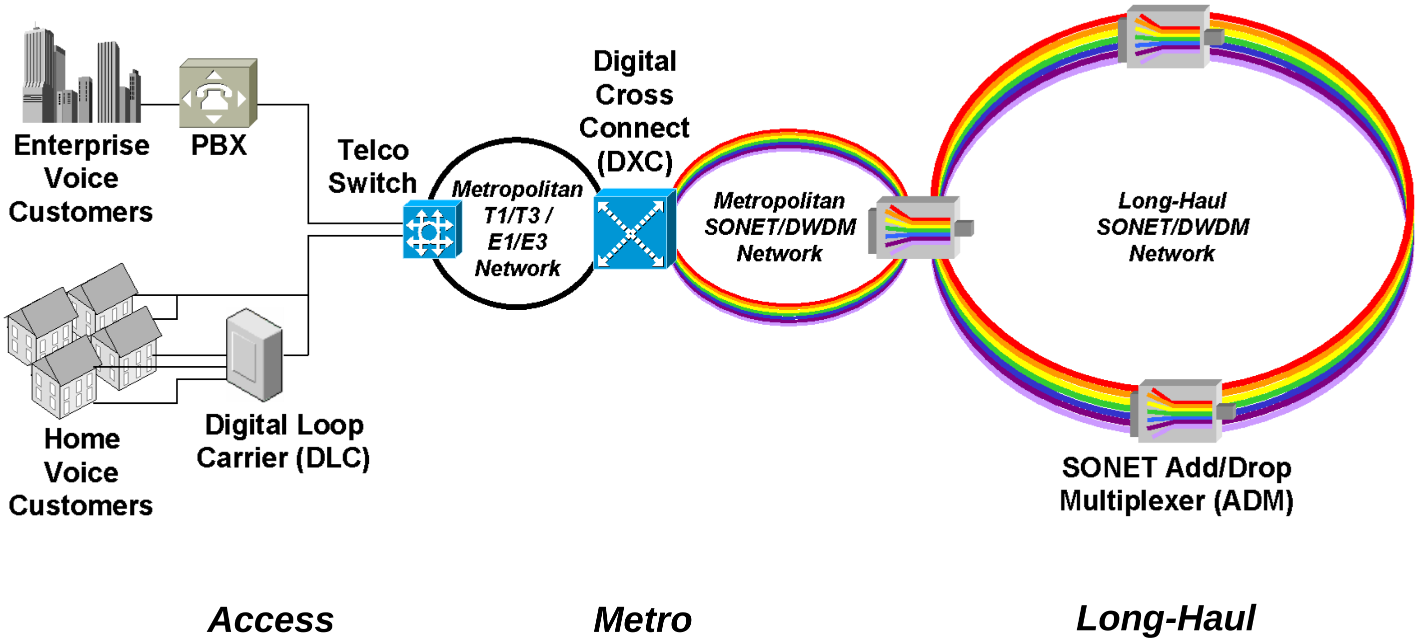

Source: Cisco Systems, Oppenheimer & Co.

The voice network was originally designed to provide circuit-switched communication between telephone users connected to a worldwide public network. It is designed for connection-oriented service, in which a user desiring service initiates a call by dialing the intended number, establishes a connection once the recipient picks up, and then begins transmitting and receiving until one party hangs up, breaking the connection. As discussed above, this is a circuit-switched network; when the caller inputs the number, the service provider determines where on the network the recipient is located and allots the appropriate amount of bandwidth (usually 64 Kbps) over each segment of the call path for the duration of the call. This ensures a consistent, uninterrupted connection.

Individual consumer telephones are connected to the public switched telephone network (PSTN) via copper twisted pair wire run from telephone poles or underground conduits directly into the home or office. These POTS (plain old telephone service) lines connect up to a line card within a telco switch housed in a carrier's central office (CO) via a dedicated line. This telco switch acts as the gateway between the individual telephone user and the PSTN.

Sometimes, multiple telephone lines will be aggregated before transport to the central office. In residential networks, a digital loop carrier somewhere in the neighborhood aggregates many lines and forwards them to the CO. In a corporate environment, an enterprise PBX (private branch exchange) handles calls between extensions within the corporate system and forwards external calls to the CO. These voice aggregation devices will usually contain POTS line cards on the front end and a T1/E1 network uplink on the back end, which will connect directly into the service provider's network.

Once the call reaches the central office, the telco switch (also called a class-5 switch or a CO switch) aggregates and switches calls. Calls headed to a local-area recipient will be connected at this juncture. Calls bound for destinations not connected to that central office are sent into the metro area network (MAN), which uses T/E lines or SONET for transport. Digital cross connects located on the MAN aggregate and switch this traffic, and calls headed to recipients in the same city will be switched and sent to the recipient's central office for call termination. These cross-connects also interface with the long-haul network, which uses higher-speed SONET transport equipment. The long-haul SONET ring will connect the call to other MANs around the country to enable long distance service.

Next-Generation Voice Network Technologies (Includes Transport)

Note that our network diagram shows a simplified representation of the traditional public voice network. A number of new technologies are being implemented to make the network faster and more efficient. Note that many of these are motivated by the desire to make voice networks more flexible when handling data and voice traffic, and thus are not pure voice or transport technologies. Some examples include:

- **Multi-Service Provisioning Platforms (MSPPs):** MSPPs enable the evolution of the optical network to support increasingly data- and voice-centric traffic streams. They combine SONET add/drop multiplexer, digital cross connect, switching, and even routing functions in a single device, and are used throughout the carrier metro network as well as the edge of the core network.
- **Optical Cross Connects (OCCs):** Optical cross connects switch optical signals without performing the optical-electrical-optical (OEO) conversion that is performed in SONET add/drop multiplexers or digital cross connects. All switching is done in the optical realm. These switches tend to be used in the core of the network, where little intelligence is required to move signals around the network.
- **VoIP for Transport:** Voice-over-IP (VoIP) brings the benefits of packet-switched networking to traditional voice telephony. With VoIP technology, calls are digitized and translated into packets, and then are addressed like data packets and sent out over the data network. Most VoIP implementations to date have been behind the scenes within service provider networks for transport, especially for international calls. Service providers simply convert standard TDM-based voice calls to IP once the call reaches their network; they then transport the packets over the data network and switch them back to TDM somewhere near the destination for call termination. This leverages the efficiency of the data network and is transparent to the customer.
- **Enterprise VoIP:** VoIP is slowly replacing the traditional PBX model in the enterprise. Instead of using traditional PSTN-based phones connected to a PBX, specialized VoIP phones connect to the Ethernet LAN, and packetized voice traffic travels alongside data packets over the LAN. A softswitch or IP PBX located somewhere in the network enables call features, and the corporate access router forwards VoIP traffic to the service provider alongside WAN-bound data traffic.
- **Residential VoIP:** Carriers are looking to offer lower cost residential voice service by implementing VoIP at the customer premise, leveraging the installed base of analog phones and copper wire but moving call transport over to lower cost IP networks. Under this model, traditional analog phones hook into a specialized VoIP gateway device at the customer premise (sometimes called a VoIP terminal adapter or VoIP CPE device) that is connected to the customer's broadband modem (which we discuss in the data network section on the next page). This gateway translates voice into IP and sends it over the broadband connection. Back in the carrier's network, a specialized voice server known as a softswitch enables call features. As in the pure-IP enterprise implementation, calls placed over the IP network bound for a recipient not on the network must be terminated by the service provider; if the recipient is on a TDM network, the call must be translated back to TDM using a media gateway in the carrier's network.
- **Cable Voice:** Cable MSOs are rapidly moving toward triple-play services, offering voice service in addition to the video and data they provide today. The long-term topology for this functionality is VoIP, with a terminal adapter or voice-enabled cable modem at the customer premise packetizing voice traffic and sending it out over the DOCSIS network. Many MSOs are rolling out this service today. In the interim, however, many MSOs have rolled out a bridge solution using circuit-switched voice; under this topology, a Network Interface Device (NID) is installed at the customer premise, which feeds traditional telephony signals over the coax network. The MSO then strips the voice signal out at the head end, and uses traditional class-5 switches and digital cross connects to enable voice service. Sometimes branded "digital voice," this service is rapidly giving way to full cable VoIP.

Telecom/Datacom

The Data Network

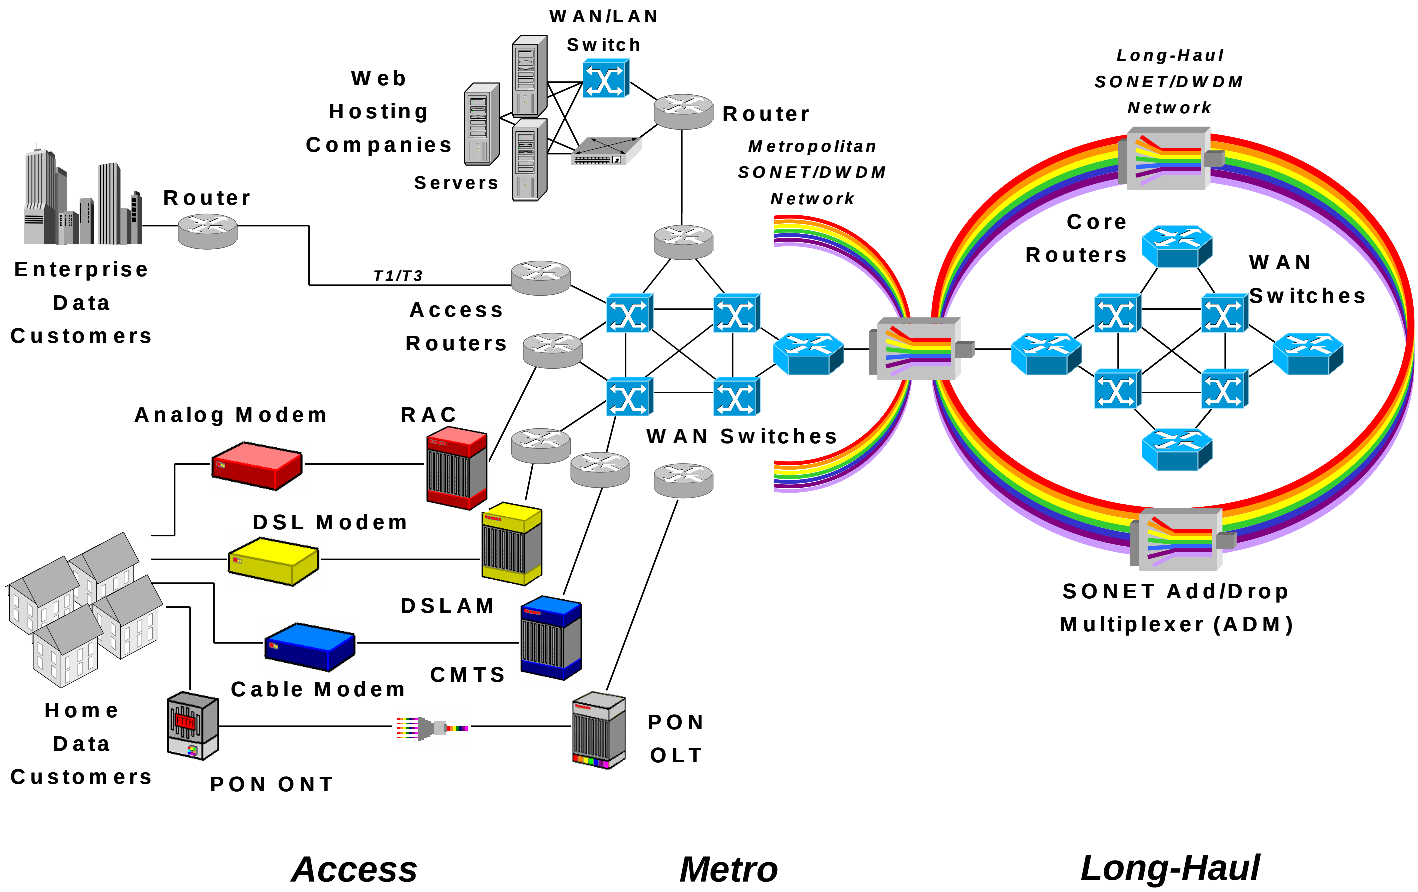

Source: Cisco Systems, Oppenheimer & Co.

The public data network uses the same physical network as the voice network does, but overlays data equipment both at the customer and at service provider points-of-presence to enable data capability. Conceptually, data equipment sits "behind" voice equipment at every point in the network, and uses the voice telecom equipment discussed in the preceding section for transport between network nodes.

In the data network, the POTS and fiber infrastructure run by the RBOCs and the coaxial cable infrastructure run by the MSOs are used by consumers to connect to the public data network. Consumers can use analog or DSL modems to connect over the POTS infrastructure, cable modems to connect over the coaxial cable infrastructure, or PON ONTs or ONUs (passive optical networking optical network terminals or units) to connect over new FTTH deployments by telcos. These lines connect into remote access concentrator, DSL access multiplexer (DSLAMs), cable modem termination system (CMTS), or PON OLT (optical line terminal) data traffic aggregators located at the central office or cable headend. Collectively, this represents the residential portion of the access segment of the data network.

On the enterprise side, corporate customers will usually have their own access links, connecting directly to the public network over a T/E line or metro Ethernet. The more sophisticated data security, reliability, and bandwidth requirements of corporate enterprises warrant these always-on connections. Note that these T/E lines can carry both voice calls and data traffic and are usually sold as a package by telco providers.

Once in the public data network, data packets are routed and switched by access routers and carrier-class WAN switches located at network POPs. These switches will typically use ATM, though newer designs use Ethernet. The transport protocol here would normally be T/E or low-speed SONET, as switching protocols like ATM and MPLS operate at layer 2 and therefore need to be encapsulated into a layer 1 technology like T/E, SONET, or Ethernet.

If the destination address is in a distant part of the WAN, packets are forwarded to the long-haul network, where core routers and larger WAN switches direct the traffic around the network. In the long-haul network, high-speed SONET is used for transport. Packets that are bound for World Wide Web addresses are forwarded to web hosting companies or corporate web servers, which connect the public data network via enterprise-class switches and routers the same way enterprise LANs do.

As discussed above, data networks are generally packet-switched, meaning that rather than establishing a dedicated pathway between sender and recipient, packets are addressed and sent out onto the network, where they are routed and switched until they reach their destinations. In a pure data network like an Ethernet LAN, both the switching and transport technologies are designed to use a packet-switched topology. This makes for a very efficient network, since resources are only allocated when they are needed (i.e., the network only "lights up" when there is data to send).

In WAN networks, however, voice-centric technologies like T/E and SONET are used for transport over long distances, and these are circuit-switched by design. This means that data packets sent over the WAN will be switched with packet-based technologies at network points of presence, but will actually be transported over a circuit-switched network. This concept is made even clearer when we consider access technologies, which by definition use dedicated lines between the user and the service provider.

Next-Generation Data Network Technologies

Note that our network diagram shows a simplified representation of the traditional public data network. A number of new technologies are being implemented to make the network faster and more efficient, but did not make it into our diagram or the discussion above. Some examples include:

- **MPLS (Multi Protocol Label Switching):** MPLS is a unified data-carrying service designed to carry packet-switched data over circuit-switched networks. It can be used to carry many different kinds of traffic, including IP packets, as well as native ATM, SONET, and Ethernet frames. MPLS dispenses with the cell-switching and signaling-protocol baggage of ATM, while preserving the traffic engineering and out-of-band control that made it an attractive technology for large-scale networks. MPLS works in conjunction with IP, and is generally considered a layer 2.5 technology.
- **Carrier Ethernet:** As carriers migrate their networks to IP, they are naturally attracted to using Ethernet as a switching mechanism due to its close technology tie with IP. However, carriers require a more robust solution than traditional enterprise-focused Ethernet equipment provides. New equipment is now coming to market specifically to address these issues, with the goal of enabling carrier-grade Ethernet. Carrier Ethernet leverages the low cost and packet-switching advantages of Ethernet while adding an extra layer of intelligence—specifically for QoS for SLA (service-level agreement) enforcement, availability and redundancy, and network security—that is required in the service provider network.
- **MSPPs:** Discussed in the voice network discussion above, multi-service provisioning platforms are designed to better handle data traffic over the voice network.
- **Metro Ethernet:** Service providers now offer Ethernet-based access lines to enterprise customers to replace their T/E lines. Not to be confused with carrier Ethernet.
- **Video Servers and IPTV Set-Top Boxes:** Next-generation IPTV services will require video servers to be installed in the metro network to deliver video content. IP set-top boxes are required at the customer premise to turn the data packets into video signals.
- **IMS (IP Multimedia Subsystem):** IMS is an attempt at creating a telephony standard infrastructure for IP services that is expected to evolve to a replacement to circuit-switched voice. The goal of IMS is to provide today's telephony services along with all the services, current and future, that can be found on the Internet. IMS uses Internet Protocols (IP and SIP) to accomplish this goal. Service providers would like to use IMS infrastructure to provide ubiquitous access to services via wireline, wireless, and cable.

Telecom/Datacom

Types of WAN Equipment - Voice Networks

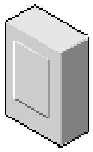

Digital Loop Carrier (DLC) - An aggregation device used in telecom networks to combine multiple telephone lines and send them over a single line to the central office. DLCs sit between the central office and customers.

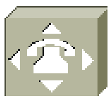

Private Branch Exchange (PBX) - An in-house telephone switching system used in enterprise networks to interconnect telephone extensions to each other and to the public switched telephone network.

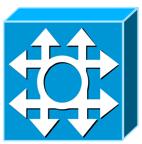

Telco Switch - Also called a “class 5 switch,” or a “CO switch,” these are the primary carrier access devices used in legacy voice networks. They are large-scale network devices containing thousands of line cards used to connect directly to voice customers.

Digital Cross Connect (DXC/DCC/DCS) - Network device used by telecom carriers and large enterprises to switch and multiplex low-speed signals onto high-speed lines and vice versa. Typically used to aggregate several T/E lines into higher speed optical SONET lines.

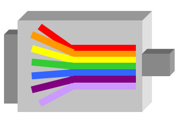

SONET Add/Drop Multiplexer (ADM) – Optical device used in fiber-based networks to enable new SONET signals to come in and existing signals to go out. ADMs are placed in intermediate positions on the SONET ring and serve as on/off ramps to other rings. Note that SONET ADMs are increasingly being replaced by Multi-Service Provisioning Platforms (MSPPs), which add digital cross connect, switching, and routing functions.

<u>**Types of Voice Networking Equipment**</u>

Digital Loop Carrier (DLC): An aggregation device used in residential telecom access networks to combine multiple telephone lines and send them to the central office. DLCs sit between the central office and customers.

Private Branch Exchange (PBX): An in-house telephone switching system used in enterprise networks to interconnect telephone extensions to each other and to the public switched telephone network.

Telco Switch: Also called a “class 5 switch” or “CO switch,” these are the primary carrier access devices used in legacy voice networks. These large-scale systems contain thousands of line cards that connect directly to customers.

Digital Cross Connect (DXC/DCC/DCS): Network device used by telecom carriers and large enterprises to switch and multiplex low-speed signals onto high-speed lines and vice versa. They are often used to aggregate several T1/T3 lines into higher speed optical SONET lines.

SONET Add/Drop Multiplexer (ADM): Optical device used in fiber-based telecom networks to enable new SONET signals to come in and existing signals to go out. As signals pass through the device, some are "dropped" by splitting them from the line, others can be "added" and directed to other destinations. ADMs are placed throughout the SONET ring and serve as on/off ramps to other rings. Note that SONET ADMs are increasingly being replaced by multi-service provisioning platforms (MSPPs), which add digital cross-connect, switching, and routing functions.

Telecom/Datacom

Types of WAN Equipment - Data Networks

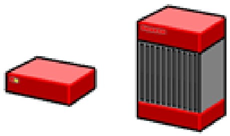

Remote Access Concentrator (RAC) & Analog Modem CPE - Equipment used to send data traffic over POTS lines in an analog format at speeds of up to 56Kbps. The RAC is housed at the central office while the modem CPE resides at the customer, embedded in a PC.

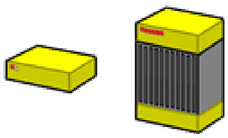

Digital Subscriber Line (DSL) Access Multiplexer (DSLAM) and DSL Modem CPE - Equipment used to send data traffic over POTS lines in a digital format at speeds of 1 Mbps or more.

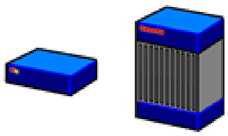

Cable Modem Termination System (CMTS) and Cable Modem CPE - Equipment used to send data traffic over coaxial cable in a digital format

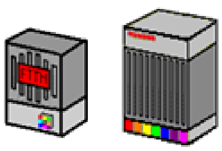

Passive Optical Networking (PON) Optical Line Terminal (OLT) and Optical Network Terminal (ONT) - Equipment used to send data over fiber-to-the-home networks at speeds of up to 10Gbps. The ONT resides at the customer with the OLT in the carrier network.

Access Router - Intermediate-speed router used at the edge of the network, both for enterprise network access and within service provider points-of-presence for routing in the metro network.

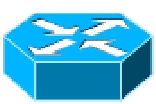

Core Router - High-speed router used in the core of the network.

WAN Switch - Network device that switches data, voice, and video traffic, used in carrier and large enterprise network backbones. Most WAN switches today use the ATM protocol and can support MPLS, though newer switches use Ethernet.

Types of Data Networking Equipment

Remote Access Concentrator (RAC) and Analog Modem CPE (Customer Premise Equipment): Access equipment used to send data traffic over POTS twisted pair copper wire in an analog format at speeds up to 56Kbps.

Digital Subscriber Line (DSL) Access Multiplexer (DSLAM) and DSL Modem CPE: Access equipment used to send DSL data traffic over the POTS copper twisted pair infrastructure in a digital format at speeds of 1 Mbps or more.

Cable Modem Termination System (CMTS) and Cable Modem CPE: Access equipment used to send DOCSIS cable modem traffic over coaxial cable video networks in a digital format

Passive Optical Networking (PON) Optical Line Terminal (OLT) and Optical Network Terminal (ONT): Equipment used to send data over fiber-to-the-home networks at speeds of up to 10Gbps. The ONT resides at the customer with the OLT in the carrier network.

Access and Core Routers: Service provider routers move traffic around the wide area data network. Access routers (1-10 Gbps) are used at the edge, both for enterprise network access and within service provider points-of-presence for routing in the metro network. Core routers (10 Gbps+) are used in the core of the network.

WAN Switch: Sometimes called a “multi-service switch,” this device switches data, voice, and video traffic in carrier and enterprise network backbones. They are multiprotocol devices, usually incorporating ATM, MPLS, and Ethernet.

Telecom/Datacom

Service Providers – Major Issues

- **Technologies driving the industry**
  - Focus remains on consolidating voice/data/video onto one network
  - Next generation access technologies – GPON in Asia and abroad; BPON/GPON in the U.S.
  - Moving intelligence from the core of the network to the edge
  - IP networks with the customers controlling more features and functionality
  - Wireless continues driving demand for increased bandwidth
- **Cable remains the biggest threat**
  - The cable companies and telcos have much to lose from aggressive price wars and both sides have some advantages
- **"Triple Play," bundling and flat rate pricing are focuses**
  - Triple play includes voice, video, and data service
  - Increases revenue per household
  - Increases the TAM over time
  - Combat OTT "chord-cutting" and required increases in bandwidth
- **Consolidation and Regulatory issues**
  - Heavy consolidation in the U.S.
  - Inter-carrier compensation – especially for VoIP
  - Universal Service Funding

The telecom industry remains in a state of transition as carriers address increasingly data-centric connectivity demands and attempt to fight off competition from cable and satellite. All service providers are moving toward triple play services, looking to provide voice, video, and data. The major issues facing carriers include:

Technology issues: Service providers are squarely focused on consolidating voice/data/video onto one network in order to lower the cost for existing services and branch into new services. This has required the adoption of next-generation access technologies, including GPON in Asia, BPON/GPON in the U.S., and next-generation DSL worldwide. Behind the scenes, carriers continue to move their networks to more flexible IP-based equipment, and are moving intelligence from the core of the network to the edges. Wireless is the largest growth driver for the industry, but the explosion of wireless data traffic has presented data congestion problems.

Threats: The biggest threat facing carriers is from cable voice, as cable MSOs prepare for broader rollouts of VoIP, and from internet calling services like Skype, Google Voice and Apple's FaceTime. Cable, offering live content streaming, may carry a competitive edge as the importance of data content and aggregation rises as consumption increasingly moves to streaming/"cord cutting." Note that the cable companies and RBOCs have much to lose from aggressive price wars and both sides have some advantages, which should prevent more aggressive behavior.

Market factors: Carriers remain focused on ARPU, in particular adding services to increase the TAM over time, but are now cognizant of the total bandwidth available for data. On the cost side, cap-ex frugality has eased, but maintenance and other variable costs remain an important focus area.

Consolidation and regulatory issues: There has been massive consolidation in the U.S. market.

Telecom/Datacom

Telecom Equipment Revenue Outlook

While the total telecom equipment revenue outlook is not the most rapidly growing of markets, there are still pockets of growth.

-Wireless infrastructure accounted for ~47% of all telecom equipment revenue in 2016 and we expect it to remain roughly flat through 2021, as a 4% CAGR in LTE spend is unable to offset lower spending across GSM, UMTS, CDMA and WiMAX.

-Routing and Switching should see above-average telecom equipment growth from 2016 to 2021 (in terms of dollars), driven by rapid growth in the datacenter/cloud optical upgrades. We forecast Routing and Switching to grow at a 3% CAGR to $29.1B by 2021.

-Enterprise WLAN continues to benefit from the rapid growth of smartphones and connected devices. We forecast the market to grow at a 6% CAGR through 2021 to $7.2B.

-Optical Transport Equipment spend is also expected to grow through 2021. We forecast a 5% C AGR from 2016 to 2021, to $20.5B.

Global Telecom Equipment Revenue

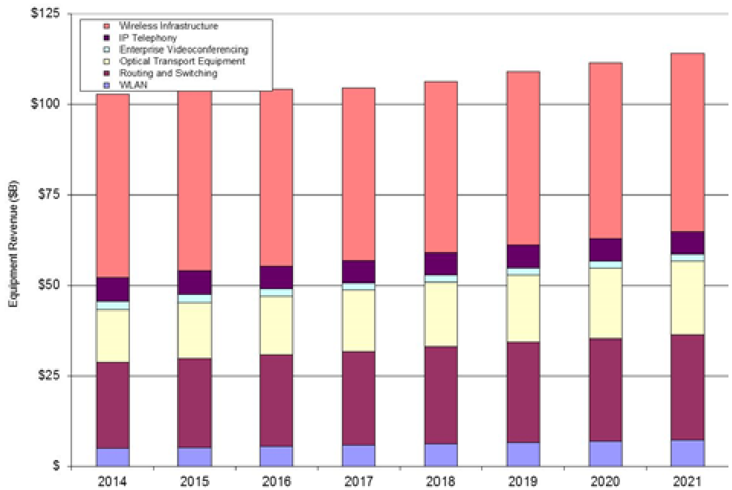

Ex-Recurring Revenue

Source: IDC, Oppenheimer & Co.

Carriers are cautious to over-spend when their own revenue growth expectations are limited, thereby limiting the total top-line growth of the telecom equipment market. However, within the nearly $100B annual market, there are pockets of growth that are likely to offset terminal declines in SONET/SDH networks, ATM switching, and WiMAX, to name a few. The fastest growing pockets of telecom equipment through 2021 will likely be in enterprise WLAN, OTN and Routing and Switching.

In total, we expect telecom infrastructure revenue, excluding recurring revenue streams, to rise from $104.2 billion in 2016 to $114.1 billion in 2021. Growth will likely be muted in North America and other developed economies, and be driven primarily by China, India and other emerging nations.

We now move on to examine in more detail three key markets within Telecom and Datacom.

Telecom/Datacom

Three Key Markets:

Modems

Data access devices used to connect individual residential customers to the service provider network over POTS copper wire or coaxial cable networks. We also include faxes due to their similarity to analog modems.

PON

Fiber access technology used to connect residential customers to the Internet using optical fiber.

Communications Infrastructure

A broad category of infrastructure equipment used for switching, aggregation, routing, and transport in voice and data networks.

We divide telecom and datacom equipment into three key segments. The first two are access technologies with transparent applications: modems and PON for data access. The third is a broad category we term communications infrastructure, which includes switching, aggregation, routing, and transport equipment used in voice and data network infrastructure.

- **Modems** (includes fax): Datacom access devices that connect individual users to the service provider network. Analog and DSL modems use the POTS telephone infrastructure; cable modems use coaxial cable networks.
- **PON:** Fiber access technology used to connect residential customers to the Internet using optical fiber.
- **Communications Infrastructure:** A broad category of equipment used for switching, aggregation, routing, and transport in voice and data networks. Note that equipment included in this category is used both in service provider networks and in enterprise networks for trunking, backhaul, and access to the public network. We chose to include all infrastructure equipment here (rather than in the networking section) due to its similarity to service provider equipment and the difficulty in distinguishing between multi-use equipment.

As in the networking section, we will be quoting a variety of speed grades throughout the datacom/telecom section that can be difficult to keep track of. The chart on the following page displays the major telecom/datacom technologies with the sub-protocols of each.

Networking Protocol Speed Chart

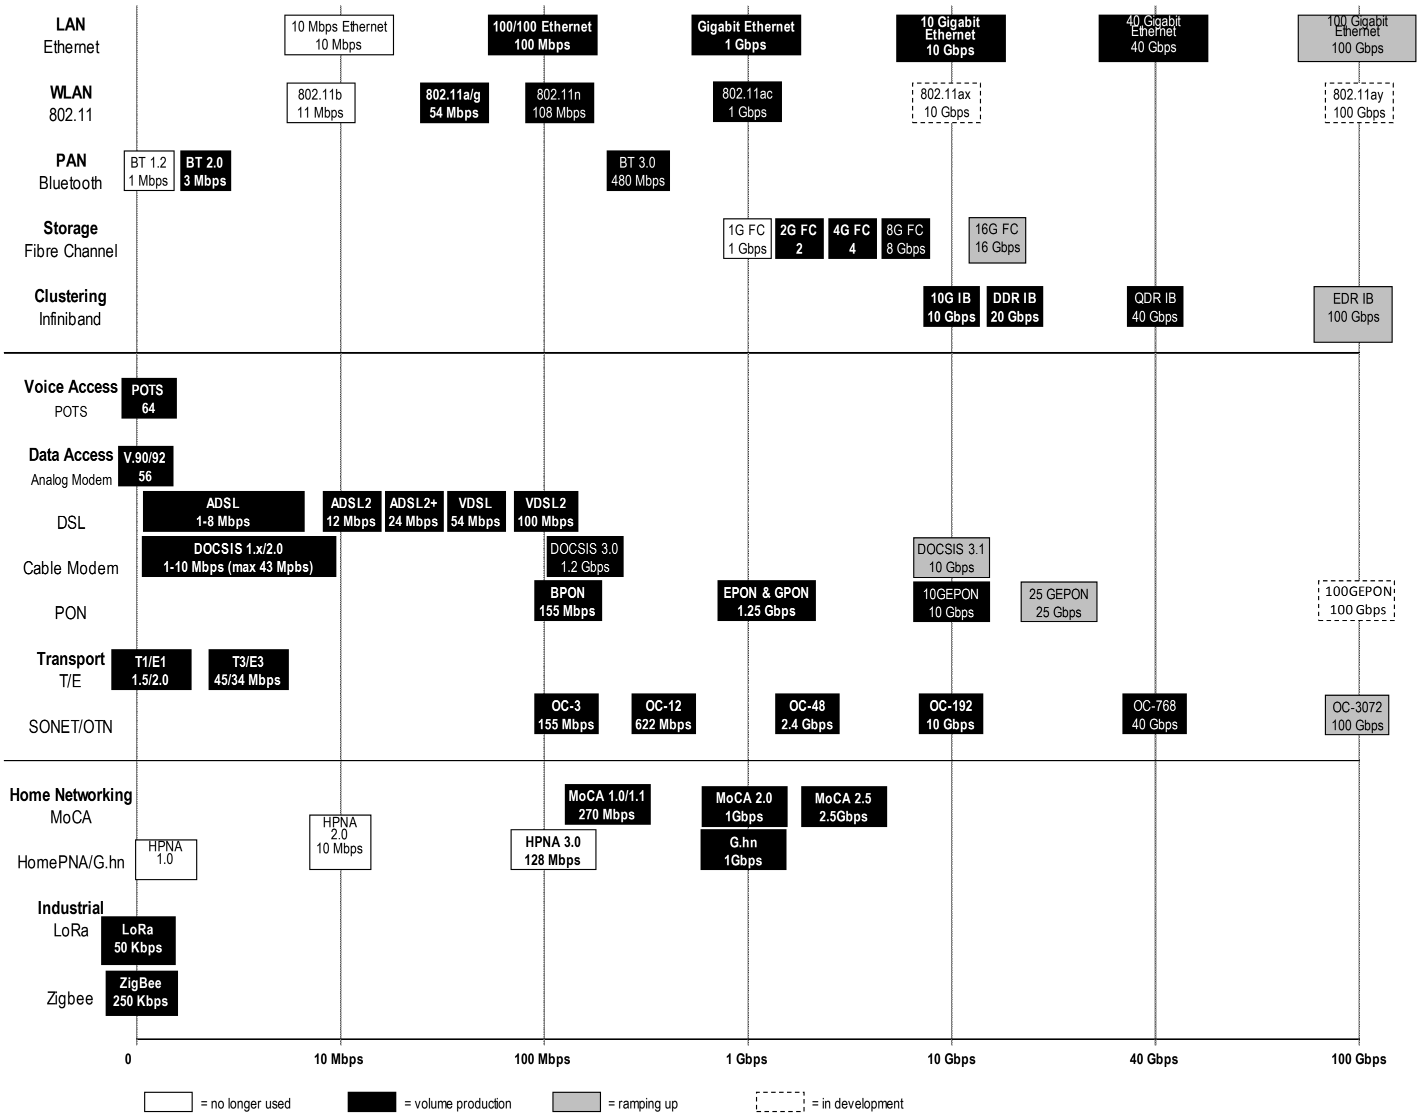

Note: In most cases, speeds quoted correspond to the maximum data rate (sometimes called the "PHY rate)" achievable according to the standard. Most networking protocols have a usable data rate that is meaningfully less than the quoted rate, though the magnitude of this difference can vary significantly depending on the protocol.
Source: Company reports and Oppenheimer & Co.

Modems

Modems are used to connect individual PCs or small consumer/SOHO networks to the Internet.

Modem CPE (consumer premise equipment) on the client side connects via copper twisted pair that has historically provided voice service or via coaxial cable that has historically provided video service. On the service provider side, a modem concentrator enables the data service and interfaces with the public data network.

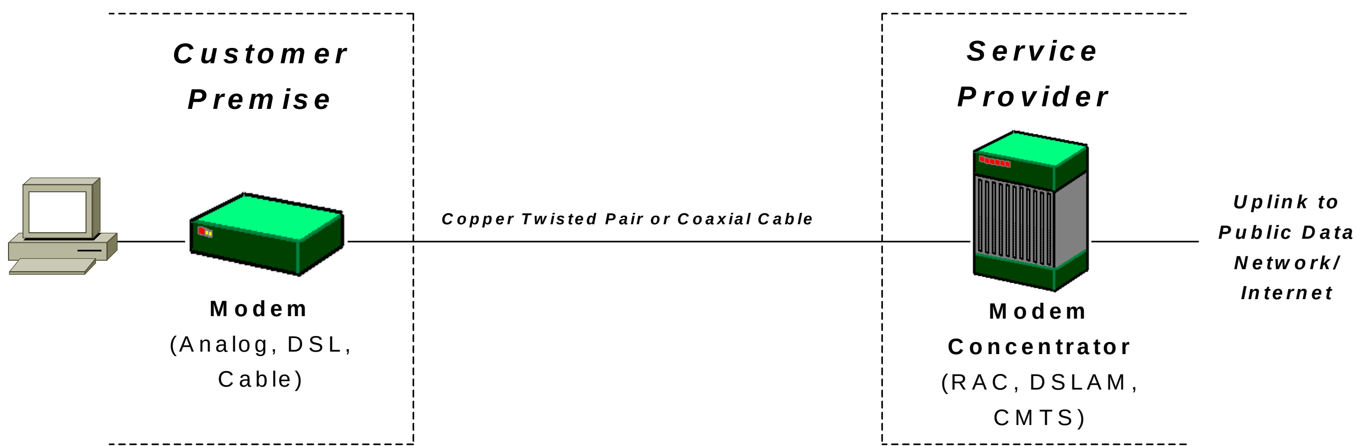

Source: Oppenheimer & Co.

Modems are the most common remote data access device, used to connect individual PCs or small consumer or SOHO networks to the public data network and the Internet. Unlike pure data networking equipment like Ethernet, modems are designed to send data traffic over networks that were originally intended for either voice or video.

Modem CPE (consumer premise equipment) on the client side connects via copper twisted pair (from the traditional voice network) or coaxial cable (from the traditional cable video network) to a modem concentrator in the service provider network. The modem concentrator acts as the gateway to the public data network. The service provider can be a traditional telecom carrier (Verizon, AT&T), a cable MSO (Comcast, Cablevision), or a specialized Internet service provider.

Modems

Three key types of modems exist:

Analog Modems - Analog modems use a traditional telephone line to place a call, which is received by the service provider modem, and transmits data in an audible analog format (like a fax machine). Transmission speed is typically 56 Kbps.

Cable Modems - Cable modems use the existing coaxial cable network to transmit data in a digital format. Transmission rates vary with user congestion, but can reach 10 Gbps downstream and upstream in new DOCSIS 3.1 modems. Note – while fiber optics are becoming more popular, fiber deployments typically end at a node before the modem, and use the existing coaxial cable network to deliver services.

DSL - Digital Subscriber Line uses the existing POTS infrastructure, but transmits data in a digital format using a higher, inaudible frequency “on top” of the traditional analog voice signal. Symmetric DSL (SDSL) transmits and receives data at the same rate, while Asymmetric DSL (ADSL) carries faster downstream rates than upstream rates. Newer, faster asymmetric protocols are referred to as VDSL (very high bit rate DSL). DSL speeds vary widely reaching up to 1Gbps

---

Three key types of modems exist:

- **Analog Modems:** Analog modems use a traditional POTS phone line to place a call, which is received by the service provider modem housed in a remote access concentrator (RAC). Analog modems transmit data in an audible analog format, which is why you can hear it if you pick up the phone line, similar to a fax machine. Top transmission speed is typically 56 Kbps. Phone lines are designed to transmit 64 Kbps of voice (the 8 Kbps difference accounts for overhead). Analog modem CPE is usually embedded on the motherboard within the PC.
- **Cable Modems:** Cable modems use the existing coaxial cable network to transmit data in a digital format. Transmission rates vary with user congestion, but can reach 10 Gbps downstream and upstream in new DOCSIS 3.1 modems. *Fiber Optics*– while fiber optics IP services are becoming more popular, fiber deployments typically end at a node before the modem, and use the existing coaxial cable network to deliver services.
- **DSL:** Digital Subscriber Line uses the existing POTS infrastructure, but transmits data in a digital format using a higher, inaudible frequency “on top” of the traditional analog voice signal. Symmetric DSL (SDSL) transmits and receives data at the same rate, while Asymmetric DSL (ADSL) carries faster downstream rates than upstream rates. Newer, faster asymmetric protocols are referred to as VDSL (very high bit rate DSL). DSL speeds vary widely reaching up to 1Gbps; however, speeds can vary dramatically depending on how far consumers reside from the central office, as noise interference can impact data rate speeds.

Modems

Cable Modem Unit Forecast

- We expect cable modem unit shipments (includes CMTS ports) to decline modestly from 45.4 million units in 2016 to 42.3 million units in 2021.
- Cable modems incorporating WLAN and VoIP should drive some incremental churn.
- Increasing trend towards chord cutting likely leads to cable modem declines
- DOCSIS 3.0 based systems are now the norm, and rollouts of DOCSIS 3.1 are now in initial stages.

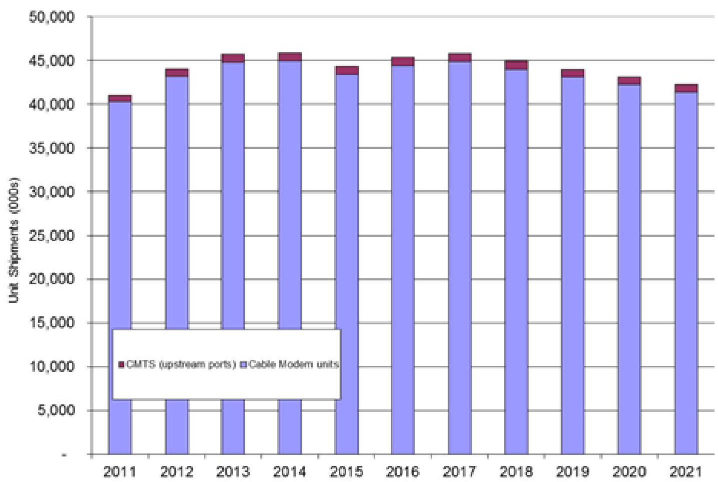

Source: Gartner, Oppenheimer & Co.

Cable Modem Suppliers

Cable Modems

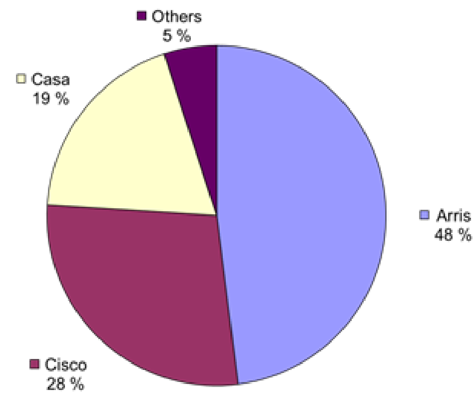

Source: IDC, Oppenheimer & Co.

Cable Modem Unit Forecast: Cable modem demand surged in the decade through 2008 as MSOs rolled out additional services and penetration neared the point of saturation. From under one million units shipped in 1998, the market reached 45.4 million units in 2016. Going forward, we expect modest decline in cable modem subscribers, as cable modem service penetration peaks and as cord cutting/over-the-top (OTT) increases. With MSOs having a stranglehold on internet services, we believe unit shipments will decline at a slower pace, as customers' appetite for over-the-top services continues to grow. In total, we expect modems to decline modestly from 45.4 million units in 2016 to 42.3 million units in 2021.

Cable Modem: In 2016, Arris led the DOCSIS 3.0 cable modem market with 48% revenue share. Cisco followed with 28%, and Casa was third with 19%.

Modems

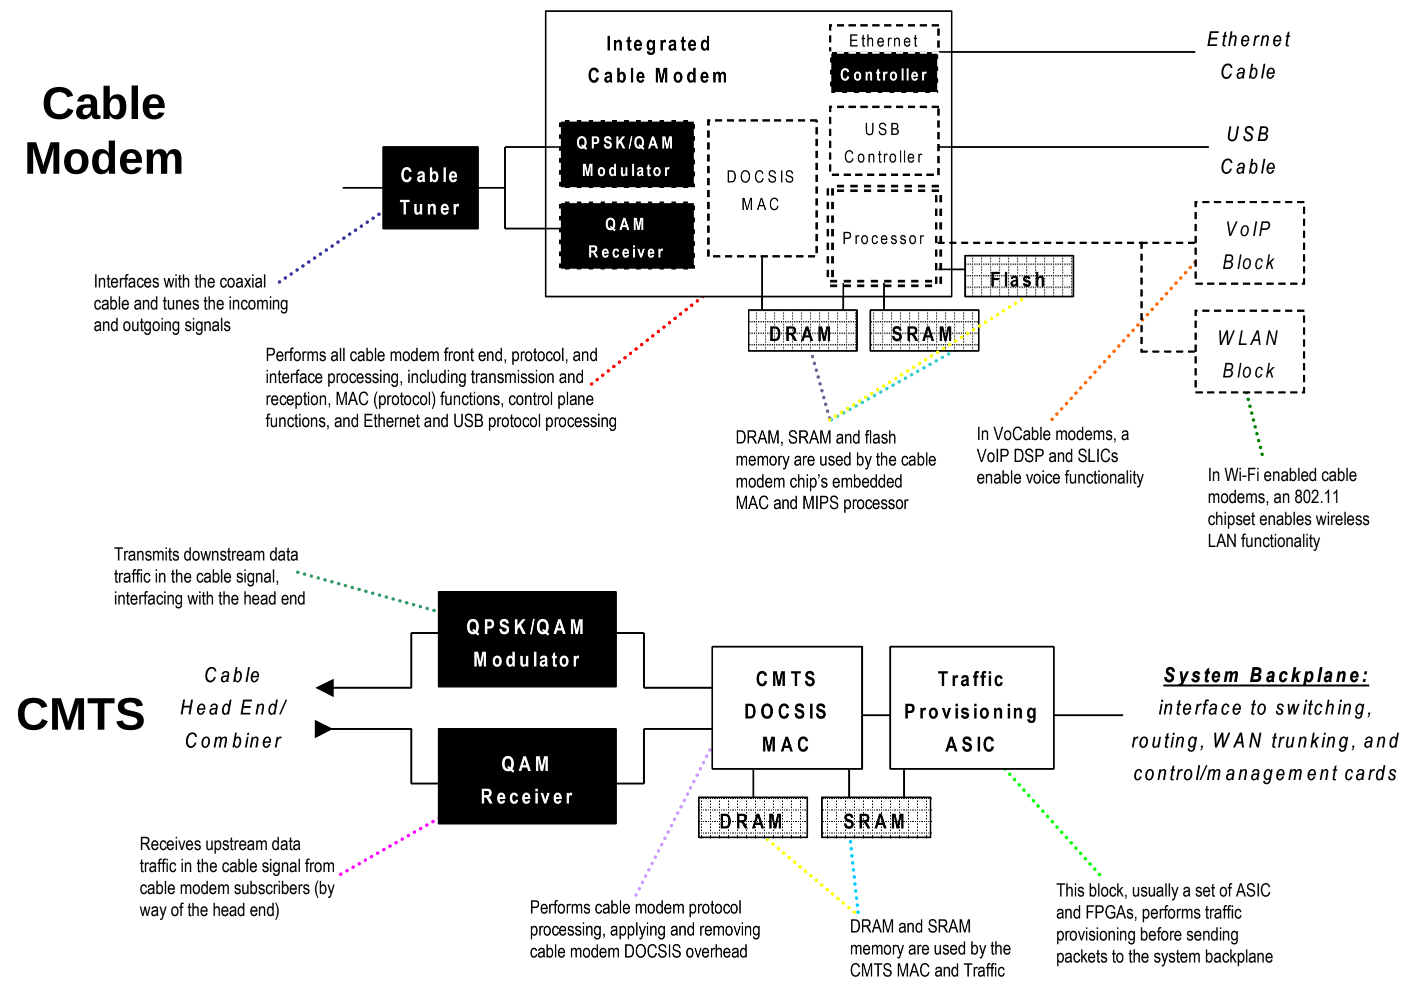

Source: Broadcom, Texas Instruments, RHK, IDC, Oppenheimer & Co.

Note: As in the analog modem section, our cable modem block diagrams, semiconductor forecast, and associated discussion include all application-specific semiconductors found in cable modem CPE and on CMTS line cards. We do not include those devices used in switching, routing, WAN trunking, and control/management cards, as these are not protocol dependent and can vary greatly depending on the functionality of the system. We discuss these ICs separately in the communications infrastructure section.

Also note that in most block diagrams in this report, we separate out IC functionality as if there were no integration to better illustrate what functions designers need in order to implement a specific application. Cable modems, however, are almost entirely based on integrated cable modem chip designs. Therefore, we picture the modem as such above.

Cable Modem CPE

- **Cable Tuner**: The tuner tunes the cable RF frequency for transmission and reception to and from the cable system. Older models use a can tuner, while newer designs use silicon tuners or a tuner block integrated into the cable modem chip.
- **Integrated Cable Modem**: The integrated cable modem includes front end processing (QPSK/QAM modulation for transmission and QAM demodulation for reception), the DOCSIS MAC (media access controller), an embedded processor and interface controllers (usually an Ethernet MAC/PHY and a USB controller). Newer cable modem chips sometimes include an integrated tuner.
- **SRAM, DRAM, flash**: Memory used by the embedded processor and DOCSIS MAC within the cable modem chip.

- **Ethernet PHY** (not shown in diagram): Some older cable modem SoC designs incorporate only the Ethernet MAC, requiring an external PHY IC to interface between the cable modem system and the Ethernet cable.
- **Ethernet Switch** (not shown in diagram): Though the basic cable modem needs only one Ethernet port to interface with the PC or wired network, newer residential gateways that include voice and wireless LAN functionality alongside the modem often include multiple Ethernet ports. These are usually implemented with a single-chip Ethernet switch, which includes Ethernet PHY, MAC, and switch functions; sometimes these functions are included in the cable modem SoC. Most gateways have 4-5 ports of 10/100 Ethernet, though this can vary.

***VoIP Block** (VoCable Modem Only)*

- **VoIP DSP** (not shown in diagram): A VoIP packet processor DSP performs the TDM-to-IP conversion on IP-based phone calls. Integrated codecs perform the analog-to-digital and digital-to-analog conversion. Note that newer cable modem SoC designs often include a VoIP DSP.
- **Line driver/SLIC** (not shown in diagram): This SLIC is a line driver that interfaces with the telephone line, sending and receiving the analog signals to and from analog telephones.

***Wireless LAN Block** (Wi-Fi Enabled Cable Modem Only)*

- **Power Amp** (not shown in diagram): The power amp sits in front of the antenna and boosts signal strength.
- **RF/IF Converter (RF)** (not shown in diagram): The RF front end controls the radio wave generation on the transmit side and receives the radio wave signal on the receive side, translating between low-frequency analog signals used by the system (specifically by the IF transceiver) to high-frequency RF signals used in radio communications. In newer implementations, the RF converter and IF are integrated in a single direct conversion transceiver, as shown in the block diagram above.
- **I/Q Modem (IF)** (not shown in diagram): The IF performs all quadrature modulation of “I” and “Q” baseband signals, converting the analog radio signal to a digital stream to be processed by the baseband and vice versa. In most implementations, the IF is integrated with the RF in a single direct conversion transceiver, as shown in the block diagram.

Cable Modem Termination System (CMTS)

- **QAM Receiver/Demodulator**: The QAM demodulator receives upstream data traffic in the cable signal from cable modem subscribers.
- **QPSK/QAM Modulator**: The QPSK/QAM modulator transmits downstream data traffic in the cable signal.
- **CMTS DOCSIS MAC**: The MAC performs protocol processing, assigning and removing DOCSIS overhead on the data packet.
- **Traffic Provisioning ASIC**: This block, usually a series of ASICs and FPGAs, performs traffic provisioning before forwarding packets to the system backplane.
- **DRAM and SRAM**: DRAM and SRAM memory is used both by the DOCSIS MAC and the traffic provisioning ASIC.

Modems

Cable Modem Semiconductor Forecast

- We expect cable modem IC revenue to increase slightly, from $1,007M in 2016 to $1,038M in 2021
- ASPs should remain relatively stable on the modem side as VoIP, WLAN and additional channel recording functions/tuners are integrated into the SoC.
- Some functionality will remain discrete to capitalize on component supplier performance advantages.

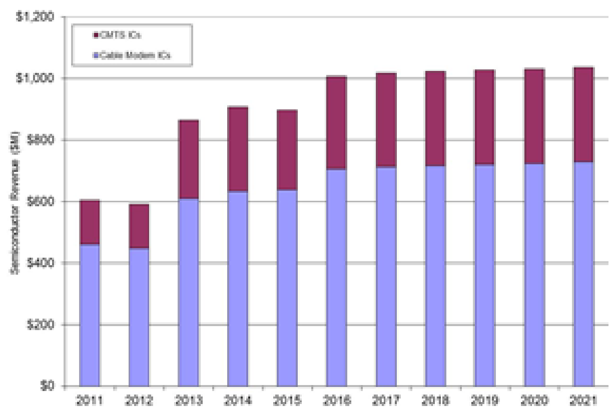

Note: Discrete WLAN and VoIP ICs used in cable modems are captured in our semiconductor forecasts for those categories.
Source: Gartner, Oppenheimer & Co.

Cable Modem Semiconductor Competitors

Broadcom
Intel
MaxLinear/Entropic

Source: Oppenheimer & Co.

Cable Modem Semiconductor Forecast: We expect cable modem ICs to increase at a modest rate over the next few years as slower growth is offset by better ASPs per box. Semiconductor revenue should increase steadily after units peak in 2017. Penetration of WLAN, VoIP and multichannel tuners/recorders are quite healthy today, and are aided by the continued deployment of DOCSIS 3.1 in developed markets. We believe that modem SoCs will continue to penetrate the market over time, gradually increasing ASPs.

Infrastructure ICs should remain roughly flat, as ASP declines and denser solutions offset growth in ports.

In total, we expect the market for cable modem semiconductors to increase slightly from $1,007 million in 2016 to $1,038 million in 2021. Modem ICs should remain around 70% of this revenue, with CMTS representing the balance.

Cable Modem CPE Semiconductor Competitors: The market for cable modem semiconductors has remained very concentrated; with Broadcom continuing to command a large share of cable modem ICs.

Modems

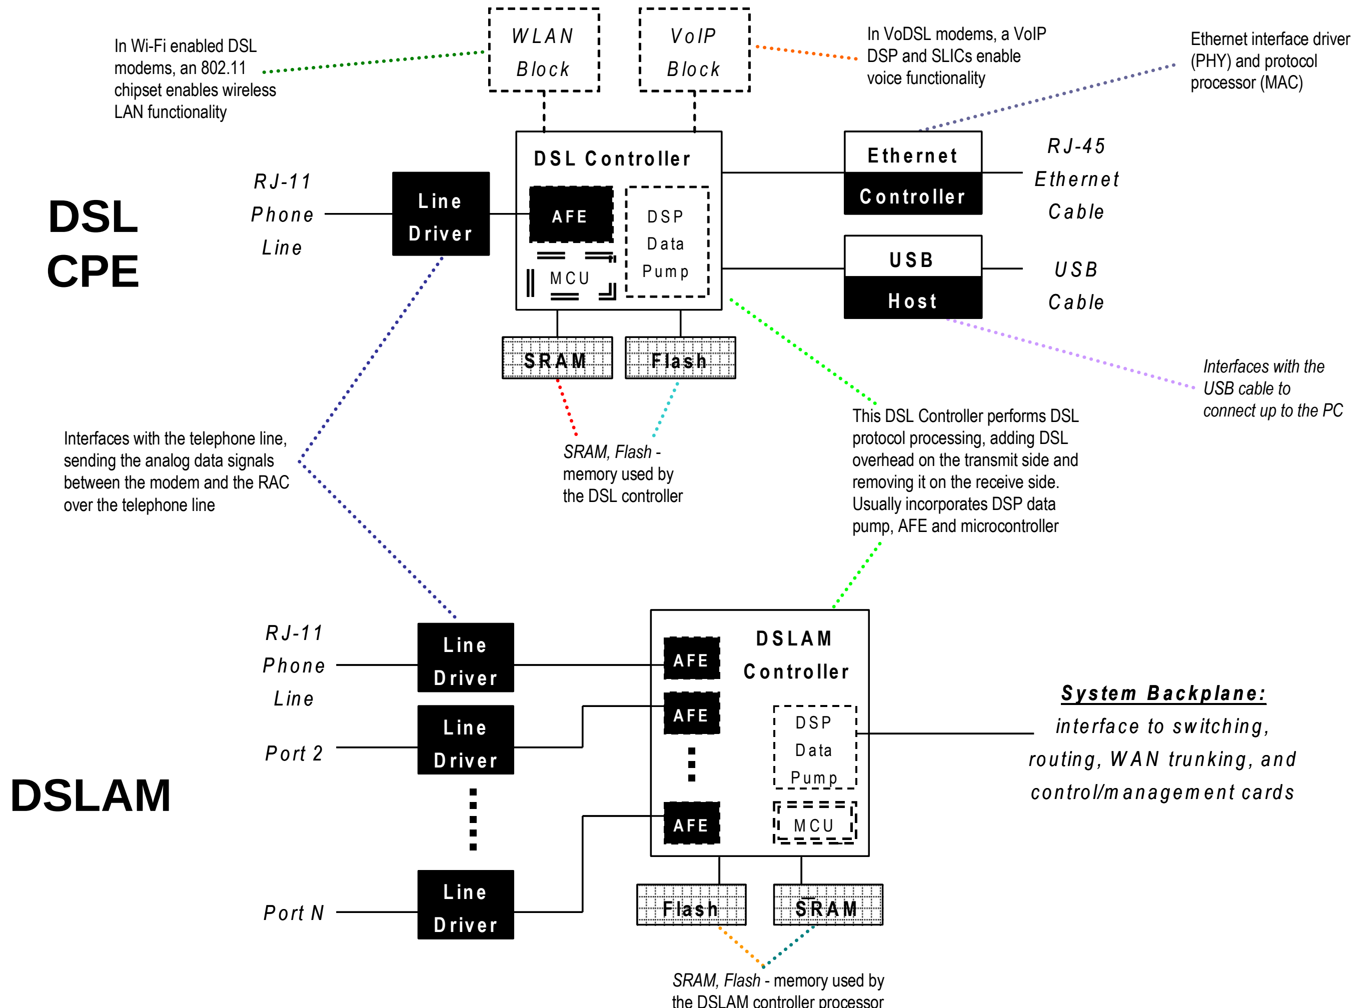

Source: IDC, Oppenheimer & Co.

Note: As in the other modem sections, our DSL block diagrams, semiconductor forecast, and associated discussion include all semiconductors included in DSL modem CPE, as well as those on DSLAM line cards. We do not include those devices used in switching, routing, WAN trunking, and control/management cards, as these are not protocol dependent and can vary greatly depending on the functionality of the system. We discuss these ICs separately in the communications infrastructure section.

- **Line Driver**: The line driver interfaces with the RJ-11 phone line.
- **DSL/DSLAM Controller**: This controller includes a DSP data pump and analog front end. It also typically includes an embedded microcontroller or microprocessor for system control. DSLAM controllers include the same functions but can include multiple AFEs and a more robust DSP.
- **SRAM and Flash:** SRAM and flash memory are used by the DSL/DSLAM controller.
- **Ethernet Controller**: An Ethernet controller provides the interface between the modem and the PC or network.
- **USB Host**: Found in USB-enabled modems, the USB host interfaces between the modem and the host PC.
- **WLAN and VoIP block:** See the cable modem diagram on for a description of the devices found in these blocks.

PON

Passive Optical Networking (PON) is a fiber-based access technology that allows service providers to deploy super-fast broadband access over optical fiber, usually referred to as Fiber-To-The-Home (FTTH). Passive splitters can be used in the field to split the optical signal into 16, 32, or even 64 channels without active switching components, dramatically lowering the cost vs. traditional metro Ethernet services.

FTTH Service Using PON

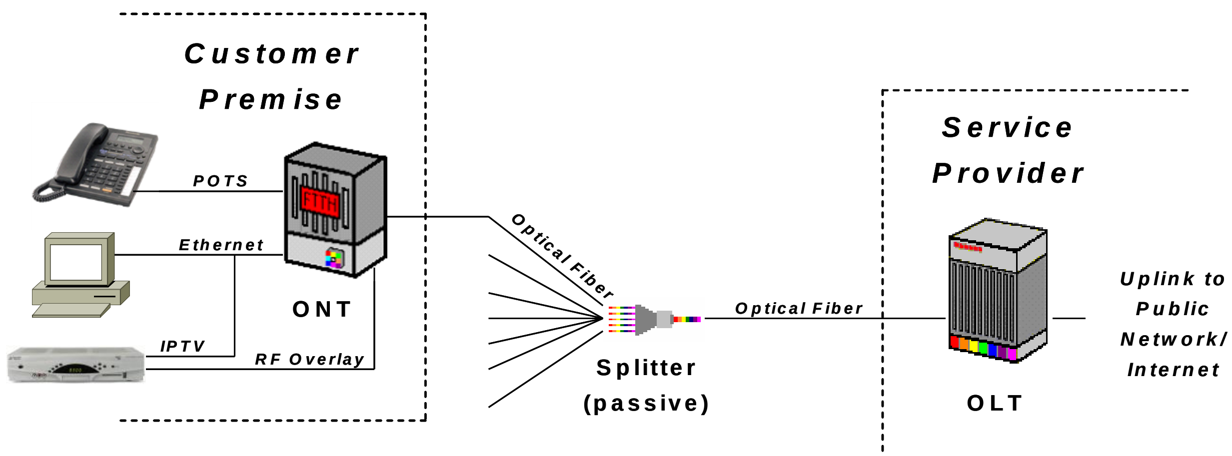

Source: Oppenheimer & Co.

Passive optical networking (PON) is a high-speed access technology that enables broadband access over optical fiber, sometimes called fiber-to-the-home (FTTH). The optical signals used in fiber-based networks are capable of achieving much higher data rates than the electrical signals that use the copper network (both the traditional POTS and coaxial cable networks). This allows for super-fast data access with a roadmap for IP-based video down the road.

In a PON system, access units called OLTs (optical light terminals) are installed in the service provider network, usually at the central office. Fiber runs from the OLT to a passive splitter out in the field, which splits the optical signal into 16, 32, or even 64 channels. Each channel is meant for an individual subscriber. Individual strands of fiber run from the splitter to customer premise units called ONTs (optical network terminals), which are installed on the exterior of the customer's home or apartment building. The ONT extracts the data (as well as TDM phone calls) from the PON signal and distributes it to the home PC or LAN over Ethernet (and POTS for the voice signals).

The key to PON technology is the use of passive splitters. These allow service providers to split off signals to individual subscribers without active switches, dramatically lowering the cost vs. traditional switched metro Ethernet services. It does, however, mean that bandwidth must be shared among all subscribers connected to the OLT.

A major driver for FTTH is the ability to provide video over the telecom network. This ideally is accomplished via IPTV, which would require an IP set-top box that is connected to the ONT. A less revolutionary option is an RF overlay; traditional RF cable signals are transmitted over the PON, and cable set-top boxes are connected to the ONT.

PON

Some service providers are hesitant to roll out full FTTH, but still want to bring fiber closer to their customers to boost data rates. In these Fiber-To-The-Curb (FTTC) or Fiber-To-The-Node (FTTN) networks, PONs extend the fiber network to within a few hundred feet of the customer, and high-speed DSL makes the final connection.

FTTC/FTTN Service Using PON

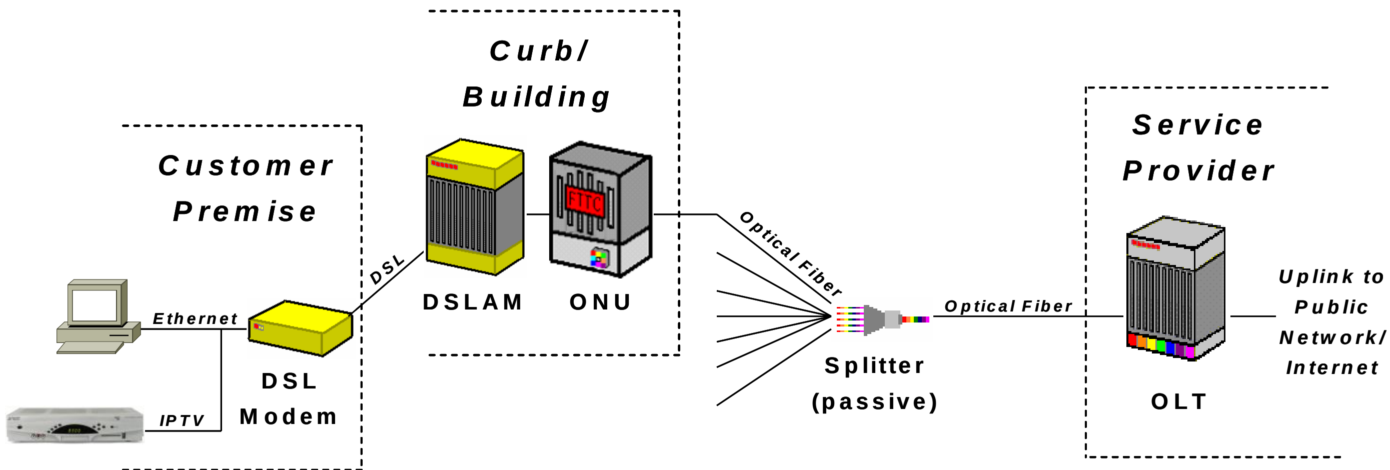

Source: Oppenheimer & Co.

Some service providers are hesitant to roll out full FTTH because of the high cost of ONTs and laying fiber around the neighborhood. They still want to bring fiber closer to their customers to boost data rates, however, and are able to leverage PON technology using a slightly different topology.

In these fiber-to-the-curb (FTTC) or fiber-to-the-node (FTTN) networks, PON technology extends the fiber network to within a few hundred feet of the customer, and high-speed DSL makes the final connection. Instead of ONTs at individual subscriber premises, the service provider installs ONUs (optical network units) in the neighborhood, at the curb, or even in the basement of a multiple dwelling unit. These ONUs interface with a small DSLAM, which distributes the signal to individual customers. Service providers are currently use VDSL2, ADSL2+, and bonded ADSL to provide this final connection, and OEMs are coming to market with equipment that merges the ONU and the DSLAM in a robust, flexible platform.

Note that FTTC/FTTN networks are capable of supporting IPTV video service, assuming the final connection has enough bandwidth to deliver the IPTV stream. They cannot, however, support the RF video overlay as in an FTTH network, since the RF overlay needs to be extracted at the customer premise.

PON

There Are Three Main PON Standards:

EPON/GE-PON – Ethernet/Gigabit Ethernet PON is based on Ethernet LAN technology. Upstream and downstream speeds can reach 1.25 Gbps and can be split into 16 channels. EPON is now being deployed aggressively in Japan and Korea; Taiwan and China are also likely to deploy EPON. There is also a move towards 10GbEPON, which can reach speeds of 10Gbps.

BPON – Broadband PON is the most common ATM-based PON technology in place today and is set by the ITU. BPON supports upstream and downstream speeds of 155 Mbps or 622 Mbps, and can be split into 32 channels. BPON can also support an RF overlay for video. BPON is seeing deployment in the U.S.; Japan had BPON service but has since moved to EPON.

GPON – Gigabit PON is the successor to BPON. GPON supports downstream speeds of 1.24 Gbps and 2.5 Gbps and upstream speeds of 155 Mbps, 622 Mbps, 1.2 Gbps, and 2.5 Gbps. 10Gbps upstream GPON is in the early stages of its rollout.

There are three main standards of PON being deployed or developed:

EPON/GE-PON – Ethernet/Gigabit Ethernet PON is based on Ethernet LAN technology and is controlled by the IEEE, the same standards body that controls Ethernet and wireless LAN. Downstream speeds can reach 1.25 Gps and use the 1500-nm wavelength. Upstream speeds can also reach 1.25 Gps and use the 1310-nm wavelength. An EPON signal can be split into 16 channels with a maximum range of around 12 miles; it does not support an RF overlay. EPON is now being deployed aggressively in Japan and Korea; Taiwan and China are also likely to deploy EPON. There is also a move toward 10GbEPON, which can reach speeds of 10Gbps.

BPON – Broadband PON is the most common ATM-based PON technology in place today and is set by the ITU. BPON supports downstream speeds of 155 Mbps or 622 Mbps, and use wavelengths of 1480-1500-nm. Upstream speeds are 155 Mbps or 622 Mbps, and use wavelengths of 1260-1360-nm. BPON can be split into 32 channels with a range of about 12 miles; it can also support an RF overlay for video using the 1550-nm wavelength. BPON is seeing deployment in the US; Japan had BPON service but has since moved to EPON.

GPON – Gigabit PON is the successor to BPON. GPON supports downstream speeds of 1.24 Gbps or 2.5 Gbps, and like BPON uses wavelengths of 1480-1500-nm. Upstream speeds are 155 Mbps, 622 Mbps, 1.24 Gbps or 2.5 Gbps, and use wavelengths of 1260-1360-nm. 10G upstream GPON is in early stages of its rollout. GPON can be split into 64 channels with a range of about 37 miles; like BPON it can support an RF overlay for video using the 1550-nm wavelength. GPON has been rolled out in multiple networks across the globe and is expected to grow at a faster rate than the overall PON market.

PON

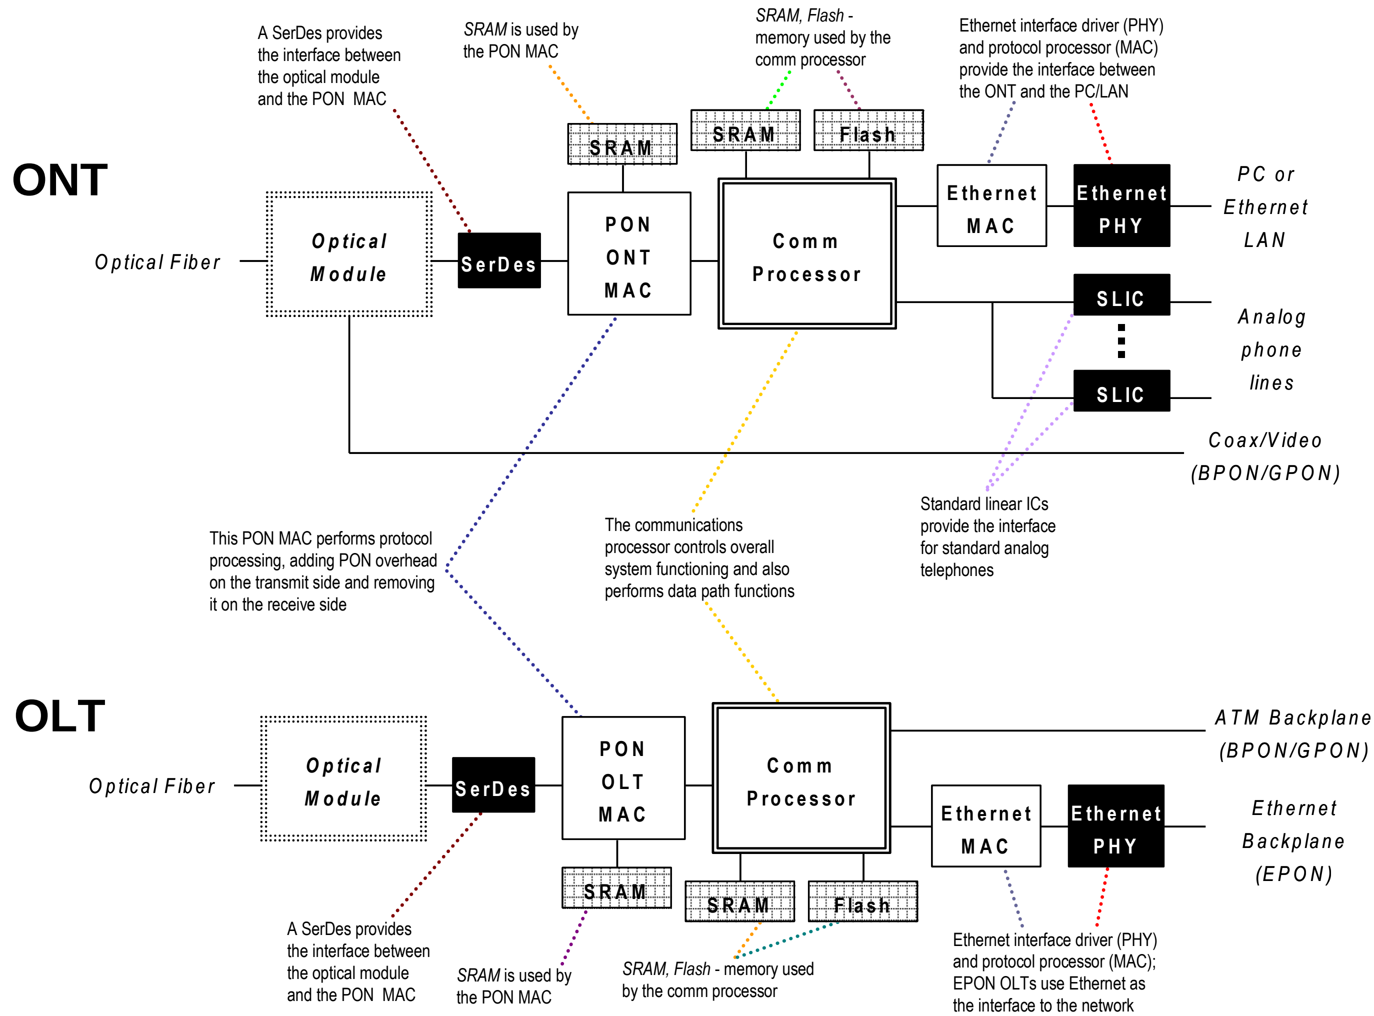

Source: Ovum-RHK, Oppenheimer & Co.

The devices found in PON equipment are not dissimilar from those in traditional broadband modems. The big exception is the optical module, which interfaces with the optical fiber.

- **PMD ICs** (inside optical module, not shown in diagram): Physical Medium Dependent (PMD) ICs like laser drivers, limiting amps, and transimpedence amps are found inside the optical module and control the interface.
- **SerDes:** The SerDes provides the interface between the optical module and the PON MAC. In newer ONT designs, the SerDes is integrated with the MAC.
- **PON MAC:** The PON MAC performs basic protocol processing, adding PON overhead on the transmit side and removing it on the receive side.
- **Comm. Processor:** The comm. processor controls overall system functioning and some data path functions.
- **SRAM and Flash:** Memory used by the communications processor; the PON MAC also uses some SRAM.
- **Ethernet MAC and PHY:** In an ONT, the Ethernet MAC and PHY provide the interface between the ONT and the host PC or Ethernet LAN. The Ethernet MAC is sometimes integrated with the PON MAC/processor. In an OLT, the Ethernet MAC and PHY feed traffic into the WAN (they are only used in EPON systems).
- **SLIC** (ONT only)**:** Standard linear ICs interface with analog telephone lines, allowing traditional telephones to be used over the PON.

PON

PON Semiconductor Market Trends

•IC vendors are targeting the U.S. market with GPON solutions.

•10GEPON deployments are gaining momentum from US cable operators and upgrades in Japan and Korea.

•MAC solutions are evolving into PON SoCs with integrated SerDes, communications processors, and memory.

•FTTC networks will use “cabinet” solutions that integrate a PON ONU with a DSLAM.

- **IC vendors are targeting the US market with GPON solutions.** The U-Verse and FiOS offerings from AT&T and Verizon saw strong growth through mid-2011, but U-Verse has slowed considerably. Verizon has gained share in the markets in which both companies compete, but has been slow to enter new markets given the significant capex requirements of laying down fiber. IC vendors are attempting to capitalize on the move to higher speed GPON. 5G and fixed wireless buildout will be another key driver to keep an eye on.
- **10GEPON deployments are gaining momentum from US cable operators and upgrades in Japan and Korea.** Driven by adoption from US cable operators, and ultimately, EPON network upgrades in Korea and Japan, the 10GEPON market will likely be the most rapidly growing PON market for the next several years.
- **MAC solutions are evolving into PON SoCs with integrated SerDes, communications processors, and memory.** Though earlier designs were built with discrete MAC, SerDes, and processor ICs, the major PON IC vendors are now in the market with SoC solutions that include these functions on a single chip.
- **FTTC/H networks will use “cabinet” solutions that integrate a PON ONU with a DSLAM.** OEMs targeting service providers that are adopting a FTTC/H (fiber to the curb/home) approach (e.g., AT&T) are developing “cabinet” systems that integrate a PON ONU with a DSLAM (ADSL2+ or VDSL/VDSL2). This should leverage the speed of PON without requiring a fiber connection all the way to the home. Typically, any system terminating within 1,000 feet of the customer is considered FTTC.

Communications Infrastructure

This section covers communications infrastructure equipment, a broad category that includes equipment used to switch, route, and manage voice and data traffic in both service provider and enterprise networks.

Included are a wide variety of dedicated infrastructure equipment, such as routers, digital cross connects, add/drop multiplexers, and multi-service switches.

We also include here the switching, routing, traffic management, and control plane functions (and their semiconductors) that are embedded in data networking, access, and wireless infrastructure devices that are also discussed in other sections of this report. Key examples include routing functions in layer 3+ switches, aggregation, routing, and WAN interface functions in modem concentrators, PON OLTs, and central office switches, and backhaul for wireless basestations.

This section covers communications infrastructure equipment, a broad category that includes equipment used to switch, aggregate, route, and manage voice and data traffic in both service provider and large enterprise networks. Included are a wide variety of dedicated infrastructure equipment, such as routers, digital cross connects, ATM switches, multi-service provisioning platforms, and optical transport equipment.

We also include here the switching, aggregation, routing, traffic management, and control plane functions that are embedded in data networking, access, and wireless infrastructure devices that are also discussed in other sections of this report. Key examples of this include routing functions in layer 3+ switches, backhaul for wireless basestations, and WAN interface functions in modem concentrators, PON OLTs, central office switches, and PBXs. One notable exception is the switching function used in Ethernet LAN equipment; we include those ICs in the Ethernet section instead of here, as LAN switching is highly specialized for that application.

Note that a unit forecast is somewhat irrelevant; we therefore display our forecast for infrastructure equipment (both for carrier and enterprise) in revenue terms. On the semiconductor side, we forecast revenue for specific ICs.. Note that in all cases, the forecast excludes all protocol-dependent ICs that are discussed in other sections; while the other sections exclude all ICs used in switching, routing, traffic management, and system control (e.g., network processors, control plane processors). Again, the one notable exception is Ethernet LAN switching, where we include switch fabric ICs in the Ethernet section instead of here.

Communications Infrastructure

The Primary Functions of Infrastructure Equipment

Though communications infrastructure equipment vary widely in application and capability, the functions they can perform are common across equipment types:

•**Transport** – moving traffic between network nodes, using copper, fiber, or coax network links external to the device.

•**Switching** – forwarding traffic between ports within a given piece of equipment.

•**Aggregation** – combining multiple smaller, slower streams into larger, faster streams.

•**Routing** – forwarding traffic by first determining the optimal path through the network.

•**Traffic Management** – prioritizing the movement of traffic based on content type, service level agreements, or other rules/algorithms.

•**Security/Access Control** – determining if a given packet or data stream should be allowed access to the network.

•**Control Plane** – management of the overall system including sub-systems listed above.

Communications infrastructure equipment can vary widely in terms of capability and application. The functions they can perform, however, are pretty common across device types. These include:

- **Transport**—moving traffic between network nodes, over network links external to the device. Most of the equipment discussed in this section uses optical fiber for transport, though copper can be used at lower speeds.
- **Switching**—forwarding traffic between ports within a given piece of equipment. This can be done either over a backplane or within a separate switching sub-system.
- **Aggregation** (also called multiplexing)—combining multiple smaller, slower streams into larger, faster streams.
- **Routing**—determining the optimal path through the network, and directing traffic accordingly.
- **Traffic Management—**prioritizing the movement of traffic based on content type, service level agreements, or other rules/algorithms. Traffic management is related to "QoS" (quality of service); infrastructure equipment can only insure a certain level of quality of service by managing the flow of traffic based on what is in the data stream.
- **Security/Access Control**—determining which packets or data streams should be allowed access to the network.
- **Control Plane—**management of the overall system, including sub-systems that perform the above-mentioned transport, switching, aggregation, routing, security, and traffic management.

Communications Infrastructure

Line Card and Chassis Systems

Infrastructure equipment is often implemented in a chassis/line card configuration. Individual cards drive transport links, and can also add switching, routing, traffic management, and system control functionality.

Telecom and datacom equipment OEMs design their carrier and large enterprise offerings to be robust, redundant, and extremely flexible. Carriers need to adjust the voice and data capabilities of their network nodes depending on local factors, and need to be able to change these capabilities as those factors change. On the enterprise side, IT managers need to have a great deal of control over the allocation of network resource among users; they also need a network that is scalable, secure, and easy to upgrade and manage. In either case, system flexibility, scalability, security, and upgradeability can be just as important as performance and general robustness of the system.

This has prompted OEMs to design their large enterprise carrier grade systems in a chassis/line card configuration instead of enclosed discrete devices like modems or access routers. Using this architecture, carriers install a chassis, which is essentially an empty metal box equipped with a backplane, power supply, and cooling system; and populate the system with line cards to suit the needs of the network node. Cards can be installed and removed at will, protecting the initial investment in the platform while allowing for future upgradeability.

Line Cards

There are four main types of line cards found in infrastructure equipment. Each line card contains a backplane interface used to communicate with the other cards within the overall system and to pass traffic back and forth.

- **I/O Line Cards:** These cards include one or more ports used to send traffic back and forth over the broader network. These can include fiber or copper ports for POTS, DSL, DOCSIS, T/E, SONET, Ethernet, PON, or any other networking technology. Some I/O line cards support shorter-range technologies like ATM UNI or Infiniband for communication with other equipment within the network node.

- **Switching Cards:** These cards include switch fabric arrays, and are used to switch traffic between the I/O cards in the chassis.
- **WAN Trunking Cards:** For systems that interface with individual users (like modem aggregators or class 5 switches), a WAN trunking card provides the high-speed uplink into the larger network. The WAN trunking card can contain routing functionality on it, though sometimes routing functions are integrated on individual line cards or conversely on a central routing card.
- **Management/Control Cards:** These cards are used to control the overall system. They can also perform special functions on the data path when there are exceptional packets. Security/access control can be performed by the management card or on a discrete security card. Traffic monitoring can be implemented in the same fashion.

Besides I/O, switching, trunking, and management, other line cards or modules can be added to perform periphery functions. Examples include LAN interface cards, test cards, and modules with optical drives embedded in them.

The type of line cards used in each system will depend on the functionality of the system and its location in the network. An add/drop multiplexer will likely have several high-speed line cards and a switch card, along with the management card. A digital cross connect will likely have a variety of line cards of varying speeds along with a set of robust switch cards, and the management card. A class 5 switch will have hundreds of individual line cards for POTS access, a switch card, a WAN trunking card for interface to the PSTN, and the management card. A DSLAM will have several multi-port DSL I/O cards, a switch card, a WAN trunking card, and the management card. A wireless base station will have several transceiver cards (essentially I/O line cards, but wireless), a WAN uplink card, a management card, and several cards that are specific to wireless (like baseband processing and power amplifier cards). A core router will have several high-speed I/O cards, several switch cards, and a robust set of management/control cards.

Communications Infrastructure

T/E Carrier

T-carrier is the first level of aggregation calls go through when being transported over the PSTN

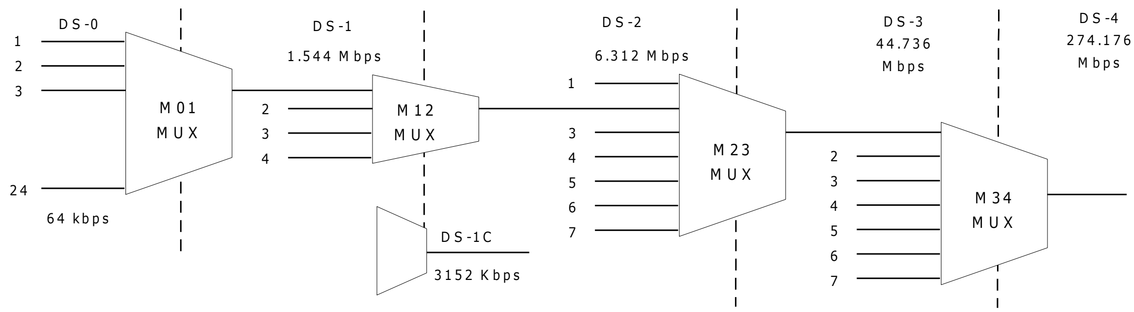

Signal Level	Number of DS-0 Channels	Transmission Rate	Geography
DS-0	1	64 kbps	North America, Europe and Japan
DS-1/T1/J1	24	1.544 Mbps	North America and Japan
E1	30	2.048 Mbps	Europe
DS-1C	48	3.152 Mbps	North America and Japan
DS-2/T2/J2	96	6.312 Mbps	North America and Japan
E2	120	8.448 Mbps	Europe
J3	480	32.064 Mbps	Japan
E3	480	34.368 Mbps	Europe
DS-3/T3	672	44.736 Mbps	North America
DS-3C	1,344	91.053 Mbps	North America
DS-3C	1,440	97.728 Mbps	Japan
E4/DS-4E	1,920	139.264 Mbps	Europe
DS-4/T4	4,032	274.176 Mbps	North America
J4	5,760	397.200 Mbps	Japan
E5	7,680	565.148 Mbps	Europe

Note: Each DS-0 channel represents a single 64 Kbps phone line
Source: Oppenheimer & Co.

T-carrier is the first level of aggregation that telephone calls or data traffic go through when being transported over the PSTN (public switched telephone network). A basic DS-1/T1 signal multiplexes 24 DS-0 (each DS-0 is a single phone call, which is a 64 Kbps signal) channels into a single 1.544 Mbps signal, called a DS-1 or T1. The next level up is DS-2/T2, which multiplexes 4 DS-1s together in a 6.312 Mbps stream (in practice, DS-2 is rarely used), followed by DS-3/T3, which multiplexes 28 DS-1 signals into a 44.736 Mbps DS-3 or T3 signal.

DS-1 and DS-3 signals are used by service providers when sending voice traffic around the metro or wide area network. Businesses also typically lease T1 (or T3) lines from the service provider that can provide a constant 1.544 Mbps (or 44.736 Mbps) channel, enough for 24 (or 672) simultaneous phone calls, which will be aggregated at the customer premise and sent as a digital signal (the “DS” in DS-1 and DS-3) to the PSTN. The bandwidth can also be used for the corporate data network in its interaction with the public WAN and the Internet.

Note: Although we mostly focus on signal rates found in North America, it is important to note that the European digital hierarchy, represented with an “E” rather than a “DS” or “T,” multiplexes a different number of DS-0 signals at each level of the hierarchy. For example, an E1 signal multiplexes 30 DS-0s instead of 24. From the E1 signal rate and higher, each subsequently higher signal is formed by multiplexing four signals from the preceding digital signal level; an E2 signal consists of four multiplexed E1 signals and an E3 signal consists of four multiplexed E2 signals.

Communications Infrastructure

SONET/SDH

SONET/SDH has historically been the primary optical transport protocol used in telecommunications

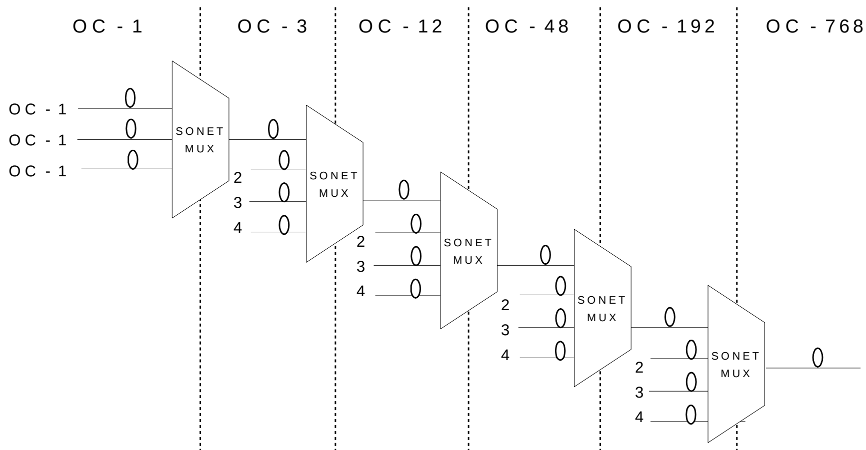

Optical Carrier Level (OC-N)	SONET STS Level (STS-N)	SDH STM Level (STM-N)	Transmission Rate	Equivalent Channels		
				DS-3	DS-1	DS-0
OC-1	STS-1	N/A	51.84 Mbps	1	28	672
OC-3	STS-3	STM-1	155.52 Mbps	3	84	2,016
OC-12	STS-12	STM-4	622.08 Mbps	12	336	8,064
OC-48	STS-48	STM-16	2488.32 Mbps	48	1,344	32,256
OC-192	STS-192	STM-64	9.953 Gbps	192	5,376	129,024
OC-768	STS-768	STM-256	39.812 Gbps	768	21,504	516,096

Note: Each DS-0 channel represents a single 64 Kbps phone line
Source: Oppenheimer & Co.

Synchronous optical networking (SONET) and synchronous digital hierarchy (SDH) have historically been the primary protocols used by carriers for optical transport (SONET is in North America, SDH in Europe). SONET can reach much higher transmission rates than T/E-carrier, as it is specifically designed to run over optical fiber. SONET equipment can handle a wide range of speeds and therefore is used throughout the service provider network. Lower-speed SONET is used to send traffic within the metropolitan area between central offices, while higher-speeds are used to send traffic between metro areas or even countries. Voice calls and data traffic are aggregated into SONET frames that travel around the network, picking up and dropping signals as calls are placed, all at high speeds.

The basic building block in SONet networks is the STS-1 frame, also called an OC-1. The frame has 51.84 Mbps of bandwidth, enough to hold an entire DS-3 signal plus the SONET overhead. When STS-1s are aggregated, they are transmitted at faster rates, such as 155 Mbps (OC-3), 622 Mbps (OC-12), 2.5 Gbps (OC-48), 10 Gbps (OC-192), or 40 Gbps (OC-768). Most metro area SONET equipment uses OC-3, OC-12, or OC-48 while long-haul uses OC-48, OC-192, or OC-768.

The SONET/SDH network, however, is being transitioned to OTN; OTN has emerged as the protocol of choice to enable the transition from voice optimized SONET/SDH network to the next generation data-optimized network. The OTN transition began in equipment for the Core and Metro and will eventually move to Edge/Access.

Communications Infrastructure

OTN

As data traffic explodes, networks are under increasing pressure to expand capacity, and OTN has emerged as the protocol of choice to enable the transition from SONET/SDH to the data-optimized next generation network.

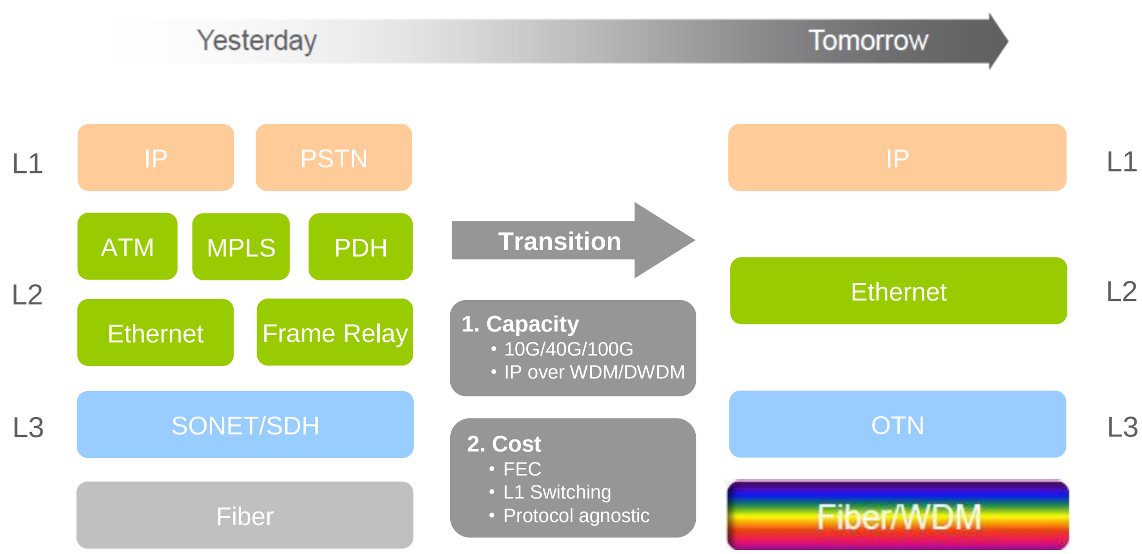

Note: Each DS-0 channel represents a single 64 Kbps phone line
Source: Oppenheimer & Co.

Optical Transport Network (OTN) is a set of optical transport standards established by the International Telecommunications Union (ITU) for the transmission of information across an optical network, officially known as ITU Standard G.709. OTN allows carriers to bridge the transition from circuit-switched, legacy transmission of voice data to Internet Protocol (IP) or packet-switched transmission of data over optical fiber. OTN, also called digital wrapper technology, has emerged as the physical layer interface protocol of choice to replace the voice optimized SONET.

In general, data traffic is better suited to packet-switched networks, which involve encapsulating data in fixed or variable length packets and sending them across a shared network. Each packet contains headers encoded with a destination address and is switched at each node in the network over the most appropriate and available physical circuit. No direct circuit or channel is established between sender and recipient as in a circuit-switched network; instead data packets bounce around in the network until they reach their destination addresses.

OTN specifications provide for a robust management overhead analogous to SONET/SDH; therefore, network operators do not have to sacrifice the ability to manage at both the payload and service levels. An OTN payload can fully encapsulate a SONET/SDH frame without terminating the SONET DCC so that Add/Drop Multiplexers (ADMs) can continue to be management in the same manner, and topology discovery will still work between customer equipment.

Communications Infrastructure

ATM (Asynchronous Transfer Mode)

ATM is a switching protocol used primarily in service provider data and multi-service networks. The topology uses switches that establish a logical circuit from end-to-end, guaranteeing quality of service (QoS), but unused bandwidth can be appropriated when needed.

ATM uses fixed-length 53-byte cells, and can switch voice, video, or data traffic at high speeds. Routers and modems feed traffic into the carrier network, which uses a series of ATM switches to handle the traffic and send it back out.

ATM is a layer 2 protocol. Cells can travel between line cards or even between switches in a single network node using point-to-point links, but need to sit inside a layer 1 protocol such as SONET or T/E to travel over the WAN.

T/E and SONET area both efficient transport protocols for the WAN, but they cannot perform the important switching functions that direct traffic around the network. Most service providers employ the switching technology ATM (asynchronous transfer mode) in the core and edge of the network in order to switch voice, data, and video traffic. Based on a cell-switching and multiplexing technology, ATM is a connection-oriented protocol for the transfer of fixed length cells, which can carry many different types of traffic. ATM equipment can specify different levels of quality of service for different cells; this is a key feature in making it the preferred choice for multi-service networks. It is highly scalable and supports transmission speeds equivalent to those of SONET.

Unlike other data link layer protocols, ATM units are not variable in length, but consist of fixed 53-byte cells. ATM switches establish a logical circuit from end to end, which guarantees quality of service. However, unlike true circuit-switched protocols that dedicate circuits end to end, unused bandwidth in ATM's logical circuits can be appropriated when needed. In this way, ATM combines the benefits of circuit switching, including constant network latency/delay and guaranteed bandwidth, with those of packet-switching, including flexibility and better network utilization. The small, fixed cell size is ideal for the transfer of voice, video and other real-time traffic, as these types of traffic are not tolerant of the long download delays associated with long, variable length packets.

Communications Infrastructure

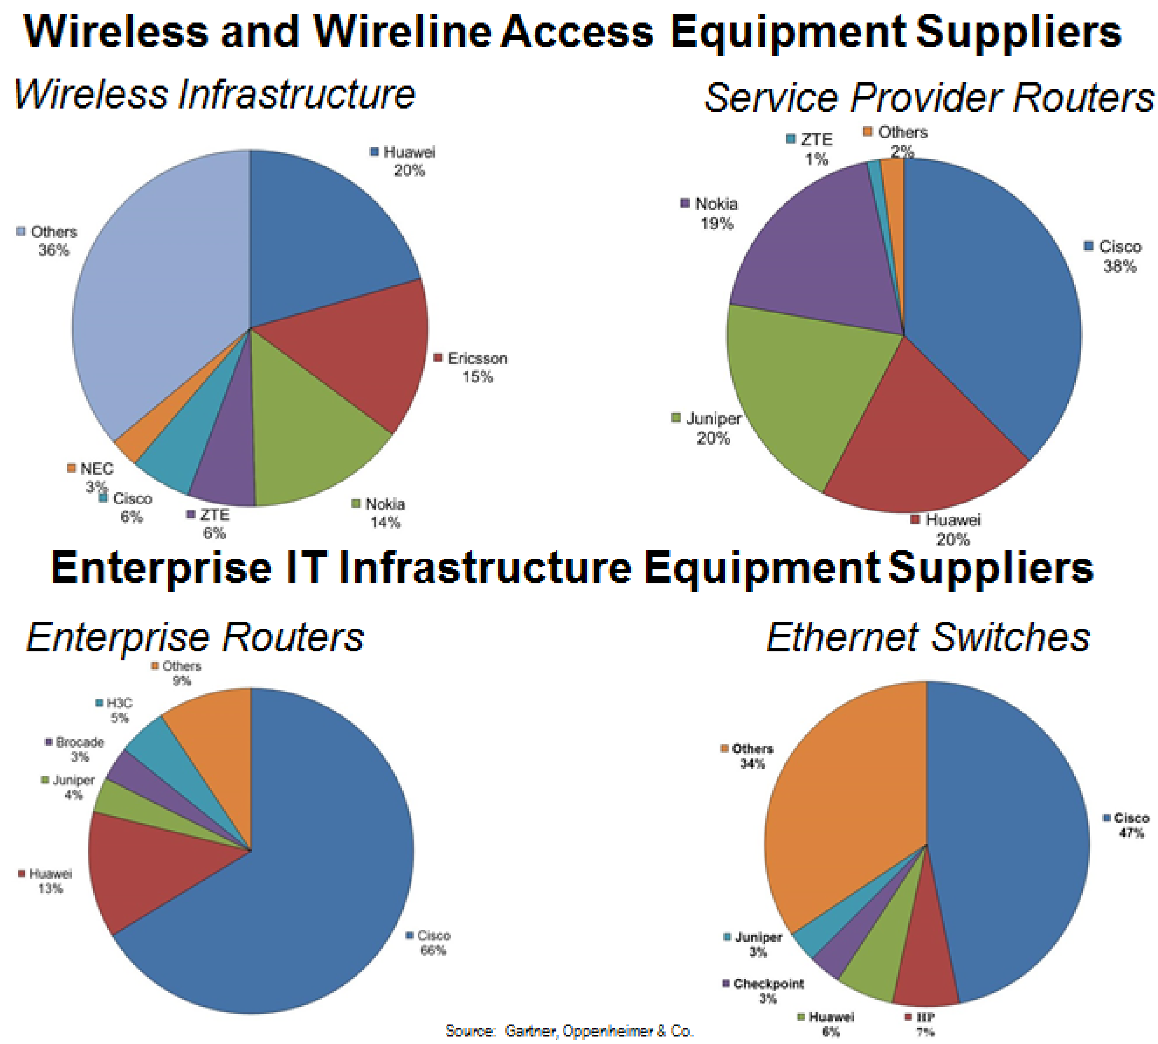

Wireless Infrastructure Suppliers: Huawei was the top supplier of wireless infrastructure equipment in 2016, with 20% revenue share. Ericsson followed with 15% share. Nokia Siemens was third with 14% share. ZTE and Cisco rounded out the top 5, each with 6% share. Another vendor was Samsung.

Service Provider Router Suppliers: Cisco was the No. 1 supplier of service provider routers in 2016 with 38% share. Huawei and Juniper were tied for No. 2 with 20% share each. Nokia held 19% market share and ZTE rounded out the top 5 with 1% share.

Enterprise Router Suppliers: The enterprise router market is dominated by Cisco with 66% revenue share in 2016. Huawei, which has gained ground in the past couple of years, was next with 13% share (down from 15% in 2014). H3C was third with 5% share while Juniper and Brocade were fourth and fifth with 4% and 3%, respectively.

Ethernet Switches: Cisco was again the dominant supplier of Ethernet switches in 2016 with 47% revenue share, while HP was a distant No. 2 at 7% share. Huawei, Checkpoint and Juniper rounded out the top players in 2016, holding 6%, 3%, and 3% market share, respectively. Other vendors include Alcatel-Lucent, Arista, Brocade, Dell and D-Link.

Communications Infrastructure

SONET/SDH Line Card

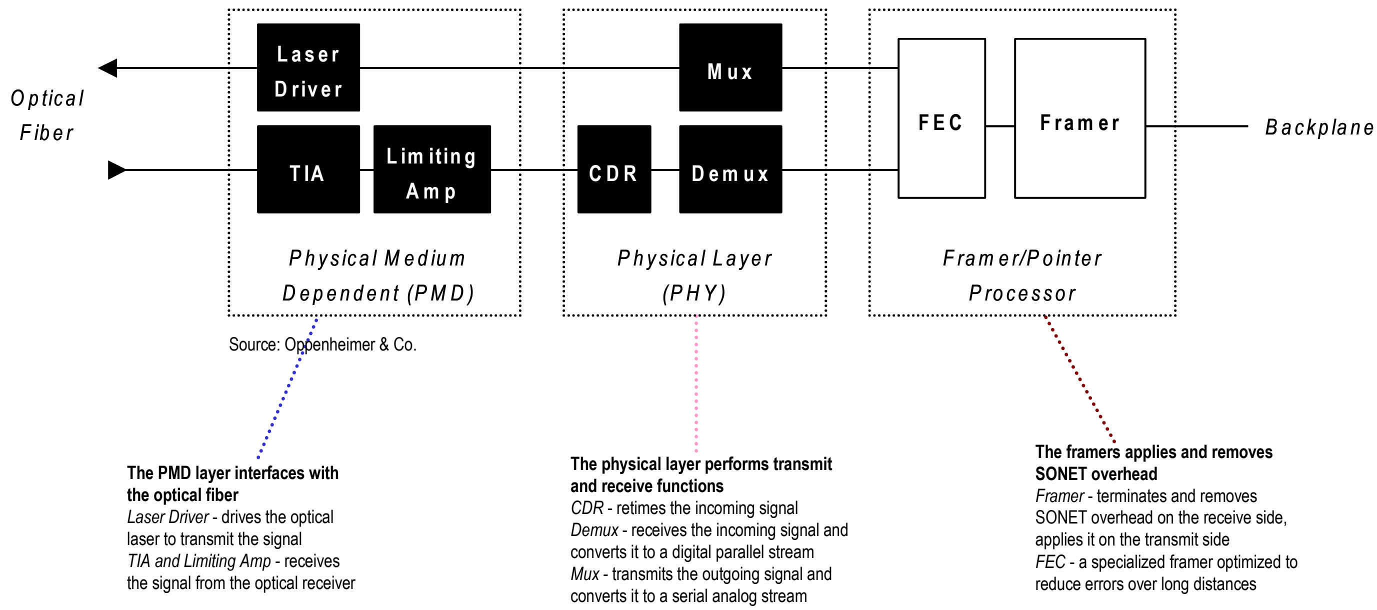

Source: Oppenheimer & Co.

SONET line cards are found in high-speed optical transport and switching systems used in metro, edge, and core networks. In a typical line card implementation, the optical signal comes off the fiber and interfaces an optical module containing a block of ICs known as the physical medium dependent (PMD) block, which controls the optical receptor and sends the analog signal into the system. Physical layer (PHY) devices sitting behind the optical module then convert the signal to a digital bit stream for processing. The framer removes SONET overhead and sends the signal to the network processing block, which classifies, modifies and manages the packet before switching. On the transmit side, the process is reversed, with the framer applying overhead, the PHY converting the signal to analog and transmitting it through the PMD devices in the optical module over the fiber.

Physical Medium Dependent (PMD)

- **Transimpedence Amplifier (TIA) and Limiting Amp**: The transimpedence amp and limiting amp receive the signal from the optical receiver.
- **Laser Driver**: On the transmit side, the laser driver controls the optical laser that sends the SONET signal across the fiber.

Physical Layer (PHY)

- **Clock and Data Recovery (CDR)**: On the receive side, the CDR retimes the signal from the PMD layer and prepares it for analog processing by the transceiver.
- **Demultiplexer (Demux)**: The demultiplexer acts as the receiver and converts an incoming signal from an analog serial bit stream to a digital parallel stream for processing by the system.
- **Multiplexer (Mux)**: The multiplexer acts as the transmitter and converts the signal from the digital parallel stream used by the system to the analog serial stream used by the optical components.
- **Transceiver** (not shown in diagram): A transceiver is an integrated physical layer device that performs clock and data recovery, mux, and demux functions on a single chip. Integrated transceivers are typically found supporting the slower, more mature speeds (OC-3 and OC-12), while faster next-generation speeds (OC-48, OC-192, OC-768, OC-3072) require discrete components to perform physical layer functions.

Framer Layer

- **Framer**: The framer terminates and removes SONET overhead on the receive side and applies it on the transmit side. The framer also breaks out the SONET stream into smaller bandwidth signals for more granularity in processing (e.g., it can break an OC-48 into four OC-12s or twelve OC-3s).
- **Mapper** (not shown in diagram): A mapper is similar to a framer but includes the capability to translate between speeds or protocols. Mappers are often used in systems that provide transport for other protocols, such as ATM or Ethernet, over a SONET link.
- **Forward Error Correction (FEC)**: Also called a digital wrapper, the FEC is a specialized framing device that uses the g.709 protocol to encapsulate the packet in an additional FEC frame for more reliable transmission over longer distances. FEC devices are only found in those systems designed to support g.709, generally higher speed long-haul equipment.

Communications Infrastructure

Communications Processors

Equipment used in communications infrastructure usually includes specialized microprocessors which perform control plane functions, both on individual line cards and on “management” or “supervisor” cards.

Usually an x86-, PowerPC-, or MIPS-based device, a communications processor sits above the data path and performs control plane functions, managing the functionality of the line card or the entire system.

Note that in addition to control plane functions, communications processors are often used to perform functions within the data path as programmed by software. Sometimes referred to as Integrated Communications Processors (ICPs), these devices are essentially fully-programmable network processors and are generally found on lower speed line cards.

Communications infrastructure equipment usually includes specialized microprocessors that perform control plane functions, both on individual line cards and on “management” or “supervisor” cards. Usually an x86-, PowerPC-, or MIPS-based device, a communications processor sits outside the data path and performs control plane functions, managing the functionality of the line card or the entire system. As they are fully programmable, communications processors are ideal for overall system control, which needs to be extremely flexible to support increasingly sophisticated OEM operating systems, a constantly evolving set of features, and varied customer preferences.

In addition to their function as the control plane processor, communications processors are often used to perform functions within the data path. Sometimes referred to as Integrated Communications Processors (ICPs), these devices are essentially fully programmable network processors (discussed in the following section) and are generally used on lower speed line cards. These ICPs often include specialty network processing logic around the processor core, making them look to some extent like SoCs, but with a robust processor at the core.

Note that this section includes standalone microprocessors used in large telecom and enterprise systems. It does not address processors used in modem CPE, wireless LAN access points, and other consumer networking devices, though the processors used in those applications are very similar to those found in infrastructure equipment. The market for standalone processors sold into CPE devices is quite small anyway, as processor cores are usually embedded into SoCs designed specifically for those markets.

Communications Infrastructure

Communication Processor Forecast

- We expect communications processors to grow at an 4% CAGR from $1.9 billion in 2016 to $2.3 billion in 2021.
- X86-, MIPS-, and ARM processors should all continue to see wide adoption in telecom and enterprise infrastructure equipment.
- Processors continue to perform many critical data path functions, and have not lost much ground to NPUs or other dedicated ASSPs.

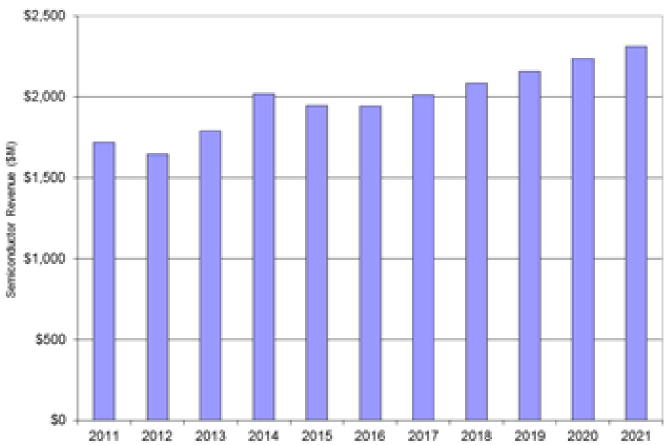

Source: IDC, Oppenheimer & Co.

Communication Processor Competitors

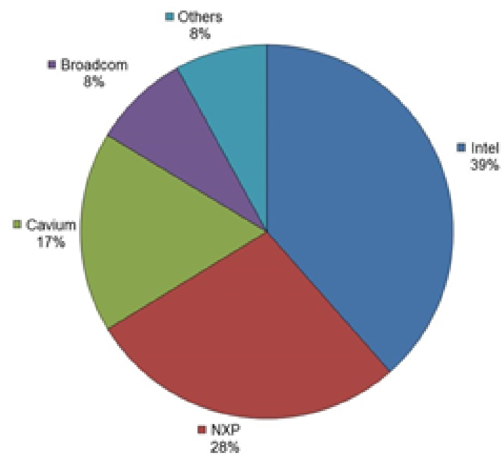

Note: Excludes processors sold into networking CPE as well as into enterprise storage equipment.
Source: IDC, Oppenheimer & Co.

Communications Processor Forecast: Communications processors are pervasive in infrastructure equipment, included on both n line cards and management cards that run the overall system. In 2016, we estimate that the market totaled $1.9 billion. Looking ahead, we expect OEMs to continue to pack additional functionality into their high-end offerings—much of which gets implemented in software—driving demand for increased processing. The trend is moving distinctly toward multicore processors. In total we expect a 4% CAGR through 2021, with the market reaching $2.3 billion.

Communications Processor Competitors: The market for communications processors is served by vendors of x86-, ARM, and MIPS-based processors. We note that both NXP/Freescale and Cavium have spent significant resources developing next-generation processors using ARM-based architecture. NXP/Freescale has long been the dominant supplier, but has lost ground in the recent past with the lack of a competitive multicore offering. Intel had 39% market share in 2016. NXP followed with 28% share. Cavium, with 17%, has gained substantial market share over the past three years by broadening beyond security accelerators to offer an extensive line of multi-core MIPS processors. Broadcom, which acquired NetLogic in 2011, held 8% share. Other players include MSCC/PMC-Sierra and MACOM. Though some of these have been aggressive in attempts to gain share vs. the leaders, most have concluded that the investment needed to compete for high-end sockets is too high, and have pulled back from this market.

Communications Infrastructure

Network Processing ICs

Communications infrastructure equipment incorporates ASICs and specialized network processing devices to perform data path processing functions. These devices primarily relate to routing and traffic management and can be included on line cards, management cards, switch cards, or WAN trunking cards.

Network Processor (NPU) - A standard programmable processor optimized for packet classification, modification, and traffic management.

Packet Classifier - A co-processor designed to classify packets based on content, usually helping to implement routing tables and QoS.

Packet Co-processor (CAM/NSE) - A specialized SRAM-based co-processor designed to quickly look up routing addresses.

Traffic Manager - A co-processor designed to manage traffic flow and implement QoS.

Security Co-processor - A co-processor designed to accelerate networking security protocols, such as IPSec.

ASICs – Custom ASICs can perform any or all of the above functions.

Communications infrastructure equipment incorporates specialized network processing devices to perform data path processing functions. These devices primarily relate to routing and traffic management and show up on control cards, switch cards, WAN trunking cards, or individual line cards, depending on the system. Key types of devices include:

- **Network Processor (NPU):** A network processor is a standard programmable processor designed specifically to perform packet classification, modification and traffic management. Software is still a major component.
- **Packet Classifier:** A packet classifier is a co-processor designed to classify packets based on content, usually helping implement routing tables and QoS.
- **Packet Co-processor (NSE/CAM):** These specialized SRAM-based devices are designed to quickly look up routing addresses. They are sometimes called “network search engines” or “Knowledge Based Processors.”
- **Traffic Manager:** A traffic manager is a co-processor designed to manage traffic flow and implement QoS. First used in ATM-based equipment, traffic management is now a common feature in most network devices.
- **Security Co-processor:** A security co-processor accelerates networking security protocols such as IPSec. These processors often occupy their own line cards within larger systems or servers to implement security protocols.
- **Network Processing ASICs:** Custom ASICs can perform one or all of the above functions.

Communications Infrastructure

Comm Infrastructure IC Market Trends

Line Card ICs

•OEMs are flocking to multi-port T/E devices at the lower speeds.

•IC vendors are increasingly integrating the T/E LIU with the framer.

•OTN has emerged as the protocol of choice to enable the transition away from legacy voice-enable SONET/SDH networks.

•ATM is being challenged by IP-based solutions.

•DSLAMs have moved from ATM to IP, basestations are also ramping.

Higher-Layer and System-Level ICs

•Communications processors are integrating advanced networking logic.

•NPUs are gaining traction, especially in access and metro applications.

•Packet co-processors continue to penetrate the LAN backbone and are being used in service provider networks as well.

•Security processors are gaining traction in routers and LAN backbones.

•IPv6 is driving the need for more advanced network processing ICs.

•Standard interconnects are gaining traction as OEMs leverage open standards like PCI-Express, RapidIO, InfiniBand, and Fibre Channel

---

SONET Semiconductor Market Trends

- **SONET PHY and framer ASICs continue to get squeezed out; the trend toward outsourcing is moving steadily up the layer stack.** Optical networking OEMs traditionally competed on performance and have therefore chosen to do much of their silicon as custom-designed ASICs. As the market has matured and as cost and time-to-market have become more important, however, OEMs have steadily moved to standard products. Today, the majority of PHY devices and framers are implemented as standard products. Higher layer devices such as network processors and switch fabrics remain primarily ASIC-based, though even this is starting to change. The availability of larger, faster and more extensively libraried FPGAs has also begun to serve the market in larger volumes.
- **Integrated SONET transceivers are becoming industry standard.** With each new speed deployment, IC vendors enter the market with discrete CDR, mux and demux chips, which they then follow up with integrated transceivers that perform all physical layer functions.
- **OTN has emerged as the protocol of choice to enable the transition away from legacy voice-enabled SONET/SDH networks.** With the birth of smartphones, mobile streaming and the explosion of the mobile-connected consumer, global data traffic has reached an inflection point; according to the Cisco Visual Networking Index, it has increased 2.4x over the past five years and will nearly double again over the next five years. Carriers have come under increasing pressure to implement capacity expansion, and OTN has emerged as the Layer 1 protocol of choice to enable the transition from voice-optimized SONET/SDH networks to a data-optimized next generation. The transition is likely to result in a substantial decline in SONET/SDH revenues and significant growth in OTN revenues over the next five years.

Communications Processor, Network Processor and Switching Market Trends

- **Communications processors are integrating more advanced networking logic, blurring the line between comm processors and NPUs, and validating the multi-core model.** Cavium has been at the forefront of pushing the multi-core transition, arguing that more cores and higher traffic capabilities could handle varied tasks more efficiently as a monolithic chip. Fellow MIPS player NetLogic (acquired by Broadcom in 2011) has followed a similar path.

- **NPUs are gaining traction, especially in access and metro applications.** Merchant solutions for network processing have not lived up to original rosy expectations, especially at the high end, where they were originally targeted. In recent years, however, IC vendors have refocused on the lower speeds, and have been successful in displacing ASICs in access and metro applications.
- **Packet co-processors continue to penetrate the LAN backbone and are being used in service provider networks as well.** Cisco has been aggressive in its use of packet co-processors in its higher-end switching and routing platforms, and has increasingly proliferated hardware-based lookup functionality across its product line. This has led to a strong growth rate for packet co-processor products over the past few years, a trend that is expected to continue through the forecast period. Besides additional customer traction, we are seeing these products move from their home in the enterprise into service provider networks, riding the trend towards IP in the carrier space.
- **Security processors are gaining traction in routers and backbone equipment.** IT security is a clear priority in enterprise networks today, and historically has been implemented with a variety of dedicated hardware and software solutions. In recent years, we have seen networking OEMs move to implement security as a feature within other network building blocks, which dramatically increases the SAM for security processors. Cisco has been particularly aggressive on this front in its pursuit of a growing TAM, as has Cavium, which has been utilizing its security and multi-core processing expertise.
- **IPv6 is driving the need for more advanced network processing ICs.** IPv6 (Internet Protocol version 6) is the most recent generation IP addressing scheme, which increases the size of IP addresses from 4 octants to 16 (which increases the total number of available addresses from 4.3 billion [or 4.3×10^{9}] with IPv4 to 340 billion billion billion billion [or 2.4×10^{38}]. IPv6 took a longer time to implement than IPv4, as it required a massive upgrade cycle of routers to support it. OEMs begun to ramp IPv6 into their product offerings in earnest in 2012. Note that on the IC side, more robust network processing ICs were required to implement IPv6; this is especially true for the compute-intensive classification and packet co-processor ICs.
- **Switch fabric IC vendors are focusing on higher capacity fabrics.** IC vendors, namely Broadcom, are designing switch fabrics with greater capacity as well as the ability to scale more efficiently. The largest capacity fabrics are predominantly ASICs today; as merchant vendors such as Broadcom ramp the capacity of their offerings we could see them encroach on the high end of the market.
- **Standard interconnects are gaining traction as OEMs leverage open standards like PCI-Express, RapidIO, InfiniBand, and Fibre Chanel.** Over time, the use of open standards should breed low cost, high-performance solutions that will also simplify system design. Note that in addition to communications infrastructure, these products find heavy usage in server (especially blade server) and storage applications.

Wireless

The cellular communications market has been the most dynamic industry of the past several years, radically transforming the way consumers communicate, connect and share. Computing power has gone mobile.

Global consumers are always on and always connected. The smartphone market has begun to mature in developed markets and has evolved into a replacement cycle. The emerging market smartphone upgrade cycle remains middle innings, shifting from feature phones to smartphones

The market for wireless communications has emerged as one of the most important in electronics. Serving both the needs of business users and consumers, wireless has changed the way people live. Wireless communications is no longer a luxury. Consumers' insatiable demand to text, chat, share, tweet, Facebook, YouTube and stream live video is driving the mobile computing revolution. Consumers are always on, always connected, and are sharing and uploading mobile content at an accelerating rate; handset OEMs, application ("apps") developers and telecom service providers are struggling to keep up.

According to the Cisco Visual Networking Index, mobile traffic grew 63% in 2016, and is projected to grow nearly 6x over the next five years. Mobile data traffic exited 2016 at 7.2 exabytes per month, up from 4.4 per month a year earlier. Roughly 429 million mobile devices and connections were added in 2016, growing to 8.0 billion in 2016. 4G traffic accounted for 69% of mobile traffic in 2016, but represented only 26% of mobile connections in 2016. 3G represented 33% of connections, and 24% of traffic. 60% of total mobile data traffic was offloaded to fixed networks vis Wi-Fi or femtocell in 2016 as well. Impressively, mobile video traffic accounted for 60% of total mobile data traffic, and mobile network speeds grew more than 3x in 2016 with average network downstream of 6.8 Mbps.

The smartphone market has become the new mobile computing platform, and content providers, consumer electronics providers, wireline telecom and cable companies, networking OEMs and application developers are rapidly evolving their business models to meet the new mobile reality. iOS and Android are the dominant operating platforms that are expected to support the explosion of mobile traffic.

The traditional computing environment has been blurred beyond recognition. The rising demands of the mobile consumer are changing the landscape for a variety of industries and traditional business models. Smartphones, and more importantly the apps which make them "smart," have changed the way consumers connect, communicate and live.

Wireless

Wireless Communications Technologies

TDMA, iDEN, GSM, PDC (2G) – The first digital cellular networks were based on time division multiplexing. TDMA and GSM have for the most part migrated to 2.5G GPRS, and PDC has gone straight to 3G W-CDMA.

CDMA/CDMA2000 1xRTT (2G/2.5G) - Pioneered by Qualcomm, CDMA technology is used by Verizon and Sprint in the U.S. and several carriers in Asia.

GPRS and **EDGE** (2.5G/2.75G) – GPRS is an update to GSM that supports data packets. Most GSM carriers are now fully GPRS-capable. EDGE is an enhancement that increases data throughput to 384 Kbps.

CDMA2000 1xEV-DO, Rev-A/B & UMB (3G/3.5G/4G) - EV-DO is the CDMA 3G implementation. 3.5G Rev-A and Rev-B boost data throughput, UMB is 4G.

W-CDMA/UMTS, TD-SCDMA and **HSPDA/HSPA+** (3G/3.5G/3.75G) – W-CDMA/UMTS is the 3G standard replacing GSM/GPRS and PDC around the world. TD-SCDMA is a variant in China. HSDPA is an enhancement for faster downlink, while HSPA+ boosts the uplink.

LTE (4G) – The "Long Term Evolution" standard is targeted at replacing W-CDMA/UMTS networks around the world. LTE uses OFDM (Orthogonal Frequency Division Multiplexing) and MIMO (multiple-input multiple-output) antenna technology.

LTE-Advanced – The current-generation LTE network supporting peak downlink data rates of 1Gbps. LTE-A is defined as "True 4G" by the ITU standards body. LTE-A is backwards compatible with LTE networks and enables Carrier Aggregation. One of the important LTE-A benefits is it is able to take advantage of optimized heterogeneous networks (a mix of macrocells, as well as picocells and femtocells).

5G - While LTE and LTE-Advanced deployments are still underway, carriers and vendors are already working on 5G technology, targeting commercialization by 2019/2020.

<u>**Wireless Communications Technologies**</u>

2G

- **TDMA**: One of the original cellular standards when wireless went from analog to digital, TDMA divides call spectrum into different time slots to transmit multiple calls simultaneously. TDMA networks have already been overlaid with the global GSM standard, and most have also been and upgraded to support GPRS and in some cases EDGE.
- **iDEN**: iDEN is a TDMA-based technology that also supports data, short messaging and dispatch radio (two-way radio).
- **GSM**: GSM is a global TDMA-based technology that uses SIM cards to identify individual subscribers. GSM is the predominant 2G standard in use throughout Europe, Asia and the Middle East and Africa, though nearly all GSM carriers have added 2.5G GPRS across the network and are adding EDGE as well. These networks will eventually migrate to W-CDMA/UMTS, and many carriers are well under way with their network upgrades.
- **PDC**: PDC was the dominant 2G standard used in Japan. Carriers in Japan have already upgraded their PDC networks to W-CDMA and are actively migrating their subscriber bases, skipping the 2.5G node entirely.
- **CDMA**: Pioneered by Qualcomm, CDMA uses a spread spectrum technique to overlap every transmission on the same carrier frequency, assigning a unique code to each signal. This requires fewer cell sites for the same number of calls. 2G CDMA networks (sometimes referred to as CDMAone) have already moved to 2.5G CDMA2000 1xRTT and are also migrating to 3G CDMA2000 1xEV-DO.

2.5G & 2.75G

- **GPRS**: GPRS is an update to GSM that supports data packets. Most GSM carriers are now fully GPRS-capable, with GPRS handset shipments now far outpacing 2G GSM shipments.
- **EDGE**: EDGE is an enhancement that increases data throughput to 384 Kbps. While GPRS is considered 2.5G, EDGE is sometimes referred to as 2.75G.
- **CDMA2000 1xRTT**: CDMA2000 is a set of standards evolving from the original CDMA technology. 1xRTT is the 2.5G version which has fully replaced 2G CDMA around the world, adding enhanced data capabilities and throughput.

3G

- **W-CDMA/UMTS**: W-CDMA is the 3G standard that is replacing GSM/GPRS (including the old TDMA networks) and PDC. W-CDMA, as the name implies, brings CDMA technology to the GSM world, and thus, it requires new network hardware in addition to different handsets (whereas GPRS was a mostly a software upgrade in the network). Carriers in Japan and Europe have already rolled out W-CDMA and are actively migrating their subscriber bases from older technologies; the US is as well but is moving more slowly.
- **CDMA2000 1xEV-DO**: EV-DO is the mainstream 3G CDMA standard that boosts throughput for the data side of the handset. The EV-DO buildout has accelerated, and EV-DO services are now available from carriers in the US and Asia. Revisions of EV-DO (known as Rev-A and Rev-B) are underway to boost data throughput and to add compatibility with WiMAX over time.
- **CDMA2000 1xEV-DV**: EV-DV was an alternate 3G CDMA standard intended to boost throughput for both the voice and data sides of the handset. Qualcomm (and therefore its handset customers and carriers) shelved plans for EV-DV chipsets and instead focused exclusively on EV-DO.
- **TD-SCDMA**: 3G standard developed from the ground up for China. The Chinese Ministry of Information Industry spearheaded the effort to drive this standard along with Chinese carriers, networking OEMs, and handset vendors.

3.5G and 3.75G

- **HSDPA**: HSDPA is an enhancement to W-CDMA/UMTS that boosts data throughput to 1.8 Mbps or 3.6 Mbps in the downlink. Further steps to 7.2 MBit/s and beyond are planned for the future.
- **HSPA+**: HSPA+ will boost the upload speed of HSDPA to 21 Mbps. While HSDPA is considered 3.5G, HSPA+ is sometimes referred to as 3.75G.
- **CDMA2000 1xEV-DO Rev-A and Rev-B**: Rev-A is a set of enhancements to EV-DO that will boost downlink data rates up to 3.1 Mbps and uplink data rates up to 1.8 Mbps in a radio channel dedicated to carrying high-speed packet data.

4G

- **LTE**: The “Long Term Evolution” standard will be targeted at replacing W-CDMA around the world. LTE uses OFDM (Orthogonal Frequency Division Multiplexing) and MIMO (multiple-input multiple-output) antenna technology that is currently used in the more advanced wireless data LANs.
    - **TDD-LTE:** Time Division Duplex LTE does not require paired spectrum and can dynamically change uplink/downlink capacity. Because it does not require paired spectrum, TD-LTE networks are a natural evolution of TD-SCDMA networks and need not be “backwards compatible” with legacy 2G/3G networks. TD-LTE will be most prominently deployed in emerging economies.
    - **FDD-LTE:** Frequency Division Duplex LTE requires different frequencies to isolate transmit/receive signals and capacity is not affected by individual uplink/downlink demands. FD-LTE is a natural migration from 2G/3G networks and requires backwards compatibility for voice signals, making it the solution for developed markets with existing infrastructure.
- **LTE-Advanced (LTE-A)**: LTE-A is the evolution of LTE into a true 4G mobile technology as set forth by the International Telecommunication Union (ITU). LTE-A must be backwards compatible with LTE technical specifications and must meet or exceed 4G specifications on peak data rates, bandwidth, peak/average spectrum efficiency, delay performance and mobility. LTE-A aims to support peak data rates of 1Gbps in downlink and 500Mbps in uplink—requiring a bandwidth of 100Mhz. Carrier Aggregation (CA) enables LTE-A to support higher peak data rates (and throughput), more efficient utilization of paired spectrum and support heterogeneous network deployments.

5G

- 5G is the next major wireless standard, which should be rolled out by 2019/2020. The Next Generation Mobile Networks Alliance defines the following requirements for 5G networks: “Data rates of tens of megabits per second should be supported for tens of thousands of users; 1 gigabit per second to be offered simultaneously to many workers on the same office floor; Several hundreds of thousands of simultaneous connections to be supported for massive sensor deployments; Spectral efficiency should be significantly enhanced compared to 4G; Coverage should be improved; Signaling efficiency should be enhanced; Latency should be reduced significantly compared to LTE. In addition to providing simply faster speeds, they predict that 5G networks also will need to meet the needs of new use cases, such as the Internet of Things as well as broadcast-like services and lifeline communication in times of natural disaster.”

Wireless

Three Key Markets:

Wireless Handsets

Wireless Infrastructure

LTE / LTE-A

Although they are intimately connected, we segment wireless communications equipment between handsets and infrastructure equipment. They are two sides of the same coin, yet the market drivers are different: handset sales are driven by consumer (and to a lesser extent enterprise) demand in the form of both new subscribers and handset replacements, and infrastructure sales are driven by carrier build-out programs and network upgrades. In fact, the design of infrastructure systems requires competencies from the wireline infrastructure world, and many of the OEM and IC vendors discussed in the telecom/datacom section also play in wireless infrastructure (whereas the handset vendors and suppliers are generally different). The design cadences and cost structure between handset and infrastructure are each also markedly different.

LTE, commonly known as 4G, is widely available in the US but is expected to remain a secular tailwind for both handset and infrastructure vendors for several years. According to Cisco, 4G connections represented only 26% of mobile connections in 2016 but accounted for 69% of mobile data traffic. 4G is not expected to represent the majority of mobile connections until 2021, when Cisco forecasts it will represent 53% of mobile connections and 79% of total traffic. By 2021, Cisco forecasts only 0.2% (25 million) connections will be on 5G, but 5G will represent 1.5% of mobile data traffic.

LTE offers substantially greater upload/download speed relative to legacy 2G/3G networks and provides consumers with a true mobile broadband experience. LTE-A is the evolution of LTE networks into a true “4G” mobile experience, supporting downlink speeds up to 1Gbps.

Wireless Handsets

Wireless handsets have seen one of the strongest adoption curves in the history of technology. Prior to taking a slight pause through the 2008/09 recession, handset shipments had experienced double-digit growth each year since 2001 until 2013. Handsets growth has decelerated since, as penetration rates have plateaued growing 6% per year from 2013-2015 before an inventory correction in 2016.

Handset Unit Shipments

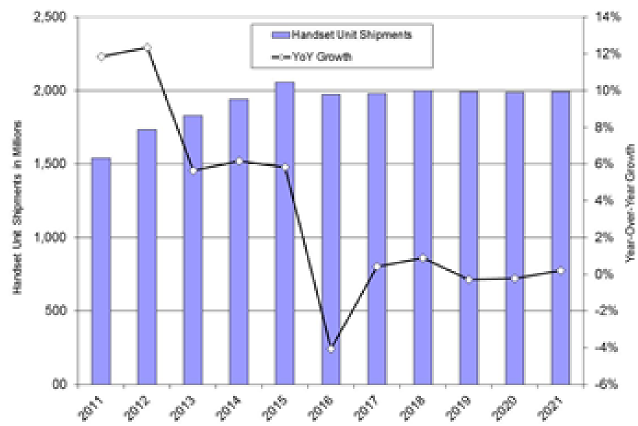

Replacement vs New Subs

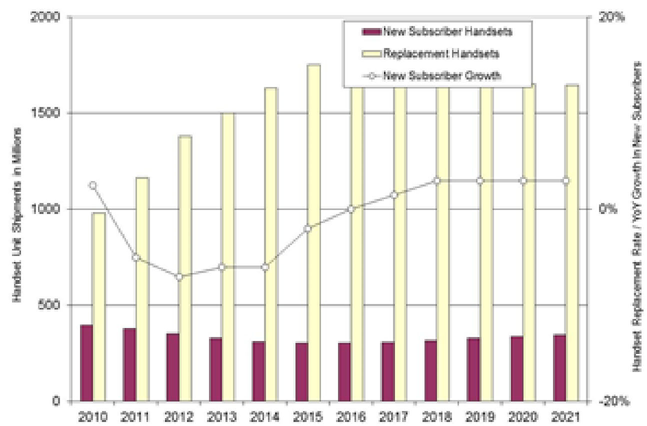

Source: IDC, Gartner, Oppenheimer & Co.

Wireless handsets have one of the strongest adoption curves in the history of technology. The late 1990s saw unfettered growth culminating in nearly 400 million handsets shipped in 2000, which seemed like a lot at the time. In 2016, handset shipments were roughly 2.0 billion.

On the subscriber handset side, two double-digit negative years in 2001 and 2002 were followed by growth rates of an average 16% in 2003-2007 and falling off by a negative 3% in 2008 by our estimates due to the worldwide recession. Just as industrialized regions were plateauing, emerging market penetration began to pick up dramatically, particularly in Asia, Latin America, and Eastern Europe. Wireless has clearly become a global phenomenon.

The largest driver of both handset unit growth and replacement rates over the next 2-3 years is likely to be global LTE rollouts and the unparalleled impact of smartphones. Developed market consumers have integrated the smartphone into the way of life and as the market matures, carriers have begun to offer incentive plans to drive replacement rates higher. Emerging market consumers are driving demand for low-cost smartphones that in many cases serve as the exclusive portal for Internet access. Smartphones are in many ways eclipsing demand for PCs and have so far transformed how the world interacts. Smartphones still represented only 45% of total mobile devices and connections in 2016, up from less than 30% in 2014 (according to Cisco's VNI Index) leaving ample headroom for growth of both the high end and low-cost variety.

Wireless Handsets

The Wireless Subscriber Base

Wireless subscribers continue to grow worldwide, led primarily by the proliferation of low-cost smartphones in emerging markets. The subscriber base in most developed markets, including the U.S. and Europe, has plateaued with population growth. As the market nears global saturation, net new subscriber growth will likely trend lower over time.

Wireless Subscribers

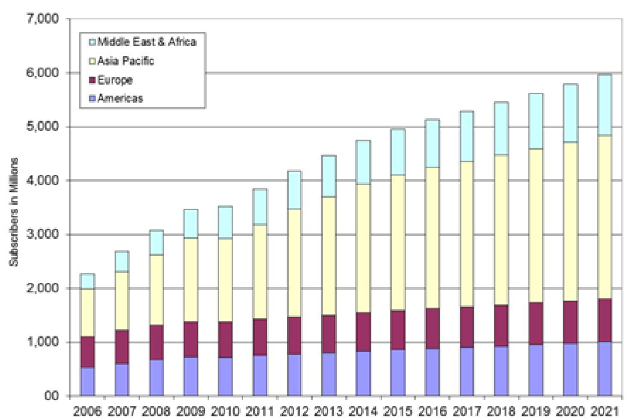

Source: Oppenheimer & Co.

New Subscriber Handsets

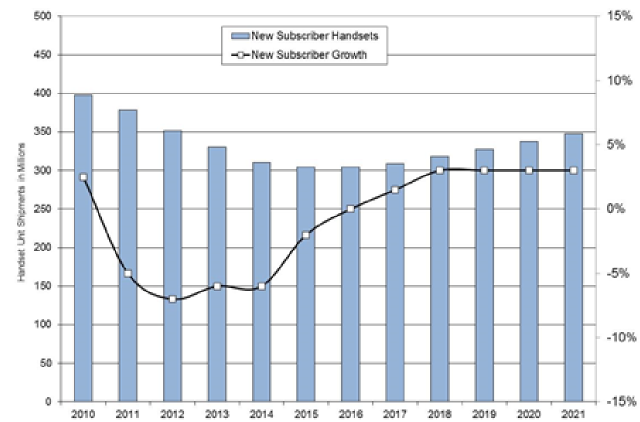

Source: Oppenheimer & Co.

After several years of 35%+ growth in the late 1990s, wireless subscriber additions trailed off in 2001-2002 (down 12% and 10%, respectively) as penetration plateaued in developed regions. The market recovered with subscriber growth increasing by an average 26% from 2003-2006 followed by 17% growth in 2007 and a 12% decrease in 2008. Subscriber growth has continued decelerating; however, with a burgeoning middle class in many emerging markets, we're forecasting acceleration of growth in 2017, with growth stabilizing around 3% annually through 2021.

Though subscriber addition growth is still healthy even after many years of robust growth, global penetration is rising, and new subscriber additions are beginning to taper off. There should still be pockets of strength in markets such as Eastern Europe, China, and India, but on the whole, net new additions should continue to exhibit slow growth in the coming years. Growth in the handset market will therefore become more dependent on replacement rates and the ongoing mix shift to smartphones, both high-end and low-cost.

In terms of the subscriber base, we find that in 2016, 17% of the world's 5.1 billion wireless subscribers were found in North and South America, 15% were in Europe, 51% in Asia-Pacific including Japan and 17% in the Middle East and Africa. APAC has risen to 51% in 2016 from just 43% in 2008. Over time, the Americas and Europe have declined as a percentage of the total while APAC and MEA have increased. Still, both the Americas and Europe regions have underpenetrated markets (particularly Latin America and Eastern Europe), which should allow the subscriber base to grow on an absolute basis.

Wireless Handsets

Wireless Handset Unit Forecast

- We expect total wireless handsets to remain roughly flat at 2.0 billion units between 2016 and 2021.
- 2G/2.5G should decline at a 14% compound rate as the 3G/4G shift continues.
- 3G/3.5G shipments should continue to decline steadily, representing 5% of total units by 2021.
- 4G is expected to aggressively grow at a 7% CAGR through 2021 and represent 84% of total shipments by this time.

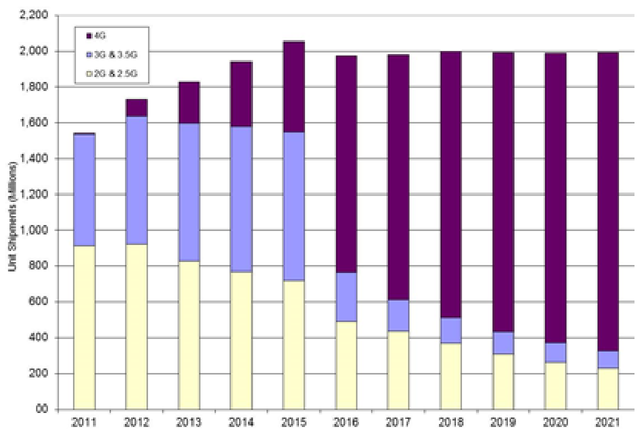

Source: IDC, Oppenheimer & Co.

Smartphone Suppliers

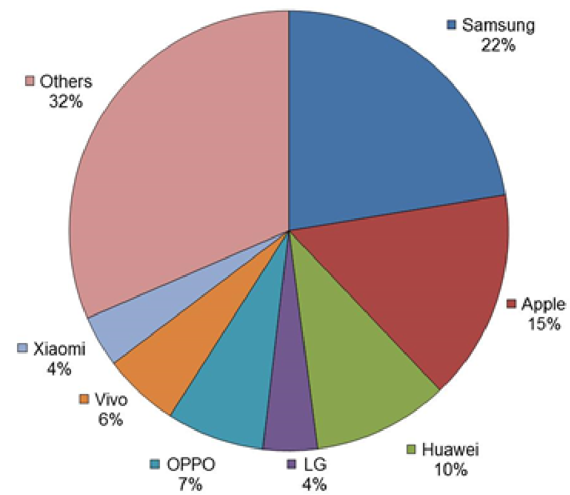

Source: IDC, Oppenheimer & Co.

Wireless Handset Unit Forecast: The handset market has seen unabated growth since 2009, led by new subscriber additions and the beginning of the smartphone upgrade cycle. Handset unit shipments in 2016 totaled 2.0 billion, down 4% year-over-year. We expect little to no growth in 2017, before modestly declining toward the end of the decade as the developed economy handset market matures and global subscriber additions naturally begin to reach saturation point.

In terms of segmentation, 2G and 2.5G handsets—which include GPRS, CDMA 1xRTT, and EDGE—represented 25% of shipments in 2016. 3G—which includes W-CDMA/HSPA, TD-SCDMA, CDMA2000/EV-DO and HSPA+—represented 14% of shipments. And 4G (LTE/LTE-A) accounted for 61% of global units.

Looking ahead, we forecast handsets to remain roughly flat at 2.0 billion between 2016 and 2021. By 2021, we expect 2G/2.5G handsets to decline to 12% of total shipments, with 3G/3.5G to represent 5% and 4G handsets to account for 84%.

Wireless Handset Suppliers: The handset market has undergone tectonic market share shifts as global smartphone adoption rises, with the emergence of low-cost Android handsets and from the explosion of "white-box" handsets in emerging markets. Smartphone share aside, the total handset market has become more fragmented. The largest supplier of handsets in 2016 was Samsung, with 22% unit share overall. Apple has grown to 15% of the total market from just 7% in 2012. Huawei was the third largest supplier with 10% share while OPPO, Vivo, LG, and Xiaomi all accounted for 7% unit share or less.

Wireless Handsets

Rising Penetration of Smartphones

- We expect smartphones to grow at a 3% CAGR, from 1.47 billion units in 2016 to 1.73 billion units in 2021, while traditional handsets will likely decline at a 10% compound rate.
- Smartphone penetration of the total handset market is expected to grow from 74% in 2016 to 85% in 2021.
- Android is still the dominant OS. We expect Android and iOS to continue to grow at the expense of all other non-iOS operating systems.

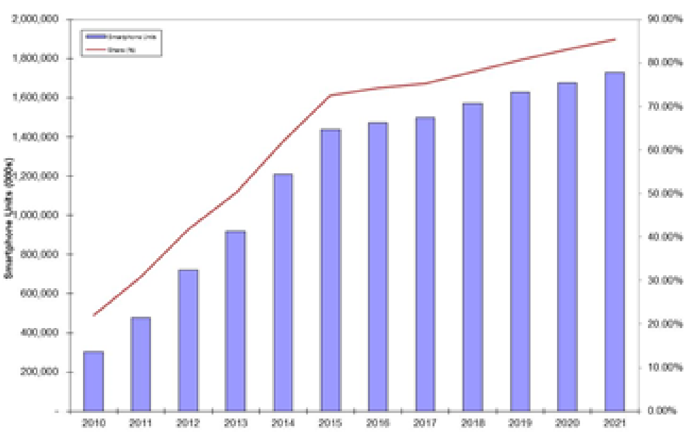

Source: IDC, Oppenheimer & Co.

Smartphone Suppliers

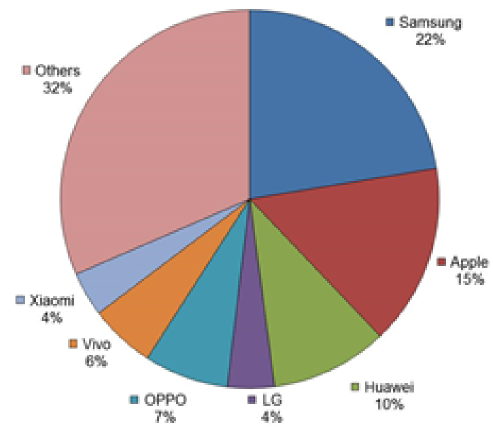

OS Share Leaders

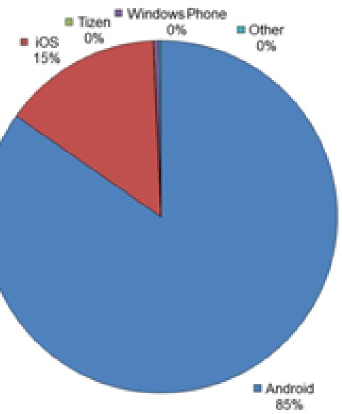

Source: 2016, IDC, Oppenheimer & Co.

Smartphone Market Share: The landscape of the handset market forever changed with the introduction of the iPhone. In an incredibly short period of time, Apple rose to market share leader before Samsung's multiple-SKU approach vaulted it to the lead with 31% share as of 1Q13. With increasing indications that the high-end of the smartphone market is maturing, we believe the most share movement volatility will be at the low-mid end of the market.

Android OS now accounts for 85% of the total market. iOS has steadily held ~15-20% share for several years. The key to Android's proliferation, we believe, has been the open nature of the platform and proliferation of the OS at both the high and low end of the market. We continue to believe a majority of smartphone unit growth going forward will come from low-cost Android platforms. While a maturing high end market and growing price competition at the OEM level are concerning trends for hardware makers, we believe the ongoing mix to 3G/4G handsets remains a sound secular tailwind for semiconductor suppliers.

Semiconductor Implications: Smartphones accounted ~45% of the total handsets in use in 2016. While high-end smartphones in developed economies have clearly reached a maturation point, the ongoing adoption of low-cost smartphones in developing economies remains very much intact. The shifts from feature phones to smartphones in emerging markets and from 3G to 4G in developing markets are sustaining semiconductor growth opportunities.

Semiconductor demands on connectivity, processing, battery life and increased complexity of the radio have increased multiple times over for high end smartphones when compared with feature phones. This presents a rising dollar content opportunity for mobile semiconductor suppliers, particularly as low-cost smartphones, with up to 5x the content of feature phones, cannibalize and ultimately replace the traditional handset market. Adding additional RF bands on LTE handsets is also a near-term boon to suppliers—all in all creating a far richer mix of semiconductor content to offset ongoing price concessions.

Wireless Handsets

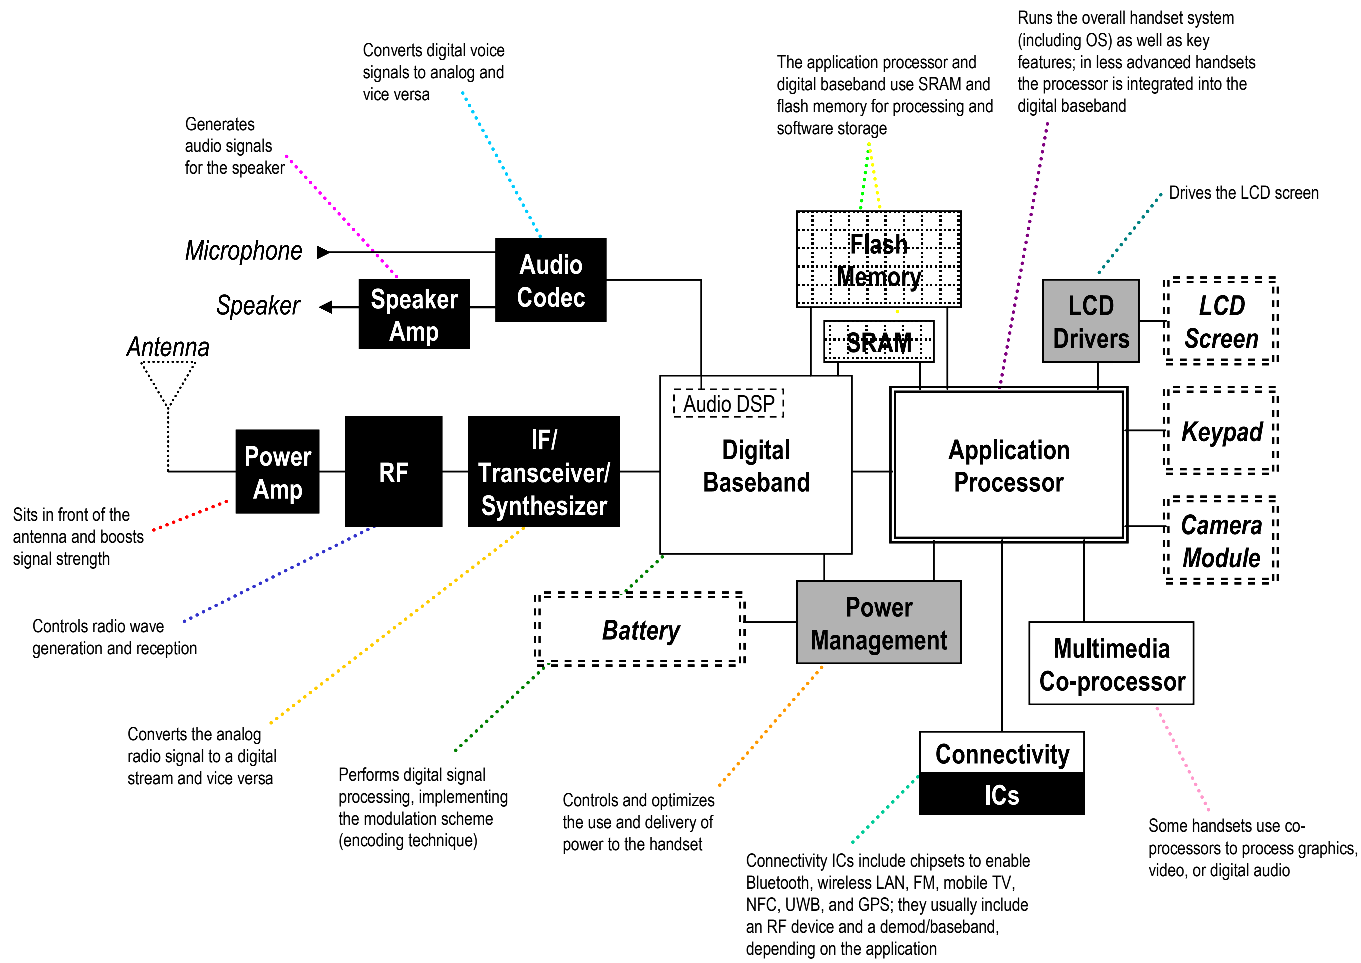

Source: Texas Instruments, Oppenheimer & Co.

The semiconductors used in wireless handsets can vary greatly by model and by OEM, especially when additional multimedia features are added, but the primary functions are essentially the same. Voice traffic is captured by the antenna and processed by the RF block, which includes a power amp, an RF and an IF transceiver, which converts the analog signal to a digital bit stream. A digital baseband processes the encoded signal into voice signals and sends it through an audio DSP to the speaker. On the transmit side, the process is reversed: the microphone sends audio signals to the audio DSP, which are processed by the baseband and sent through the RF block to the antenna.

In addition to these devices, an application processor controls the system and runs the phone operating system and other software; this processor is integrated into the digital baseband in mainstream phones. More advanced handsets will also contain discrete ICs or chipsets to enable features like the camera, graphics, Bluetooth, wireless LAN, GPS, FM, NFC, or mobile TV. The system also uses significant flash memory, SRAM and power management devices.

RF Block

- **Power Amp**: The power amp sits in front of the antenna and boosts signal strength.
- **RF**: The RF front end controls the radio wave generation on the transmit side and receives the radio wave signal on the receive side, translating between low-frequency analog signals used in the handset (specifically by the IF transceiver) and high-frequency RF signals used in radio communications.

- **IF/Transceiver/Synthesizer**: The IF performs all quadrature modulation of "I" and "Q" baseband signals, converting the analog radio signal to a digital stream to be processed by the baseband and vice versa.

Baseband and Audio Processing, Power Management

- **Digital Baseband**: The baseband processor performs digital processing on the signal, implementing the wireless modulation scheme (the encoding technique, e.g., GSM, GPRS, EDGE, CDMA, CDMA2000 etc.), encoding and extracting voice signals in and out of the transmission protocol. It is usually implemented as a DSP or an ASSP with embedded DSP.
- **Audio DSP**: The audio DSP is integrated into the digital baseband and processes the audio signals, converting the stream into voice signals and vice versa.
- **Audio Codec**: On the receive side, the audio codec converts the digital voice signals to analog signals to be sent to the speaker amp; on the transmit side, it receives and converts voice signals from microphone to digital for processing by the audio codec.
- **Speaker Amp**: The speaker amp drives the audio signals in the handset speaker or headset.
- **Power Management**: Power management ICs regulate, control, manipulate and optimize the use and delivery of battery power to handset components.
- **Analog baseband** (not shown in diagram): The audio DSP, codec, and speaker amp are often integrated into a single IC known as the analog baseband. Many analog basebands also include power management.

Application Processor and Memory

- **Application Processor**: The application processor runs the phone operating system and other software and controls overall handset functionality. Most mainstream handsets integrate the processor into the digital baseband; more advanced handsets use a discrete application processor.
- **Flash Memory**: Flash memory is the primary memory used by the handset for storage of system software, phone book entries and other user data. In traditional handset designs the flash memory used is high-density NOR flash; however this has largely been replaced by NAND.
- **Mobile DRAM:** The digital baseband and application processor use some volatile memory for processing functions. In traditional handset designs this has primarily been SRAM, though today it is largely comprised of mobile DRAM.
- **Multimedia and Connectivity ICs**: These ICs or chipsets are used in advanced phones to implement multimedia and connectivity functions. These features will usually be implemented with a chipset that includes digital processors plus an analog front end as described below. Note that some of this logic can be loaded onto the application processor or an intelligent co-processor, eliminating the need for discrete ICs.
    - **Camera:** Camera modules usually include a CMOS or CCD image sensor and an image processing IC.
    - **Advanced Graphics:** Some handsets include a graphics co-processor, though many application processors can also perform this function.
    - **Bluetooth:** Bluetooth functionality requires an additional radio and baseband, usually implemented as an integrated chip with Wireless LAN.
    - **Wireless LAN**: 802.11 functionality requires an additional radio and baseband.
    - **GPS:** GPS functionality requires an additional radio and baseband.
    - **Near Field Communications (NFC)**: Mobile payment system is now available in most high-end smartphones.

Wireless Handsets

Handset Semiconductor Forecast

•Handset ICs are expected to grow at a 5% CAGR from $80.9 billion in 2016 to $104.8 billion in 2021.
•Both discrete basebands and integrated baseband/apps processor should increase at a 3% CAGR.
•RF/power amps should rise 3% driven by rising band count and antenna complexity on LTE/LTE-A handsets.
•Connectivity ICs should decline at a 2% compound rate, while memory and power management Ics should grow at a 10% and 3% CAGR, respectively.

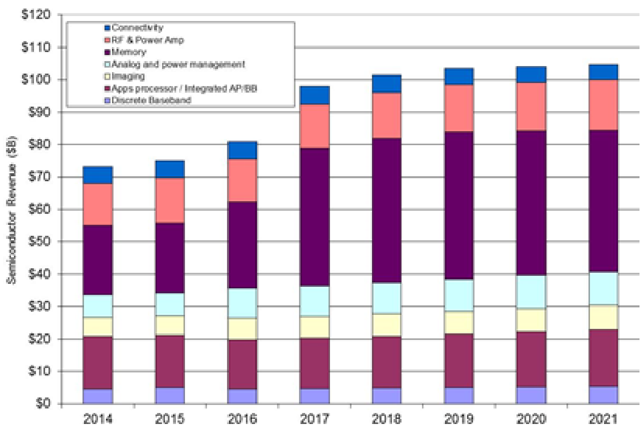

Note: Includes memory.
Source: IDC, Oppenheimer & Co.

Handset Semiconductor Competitors

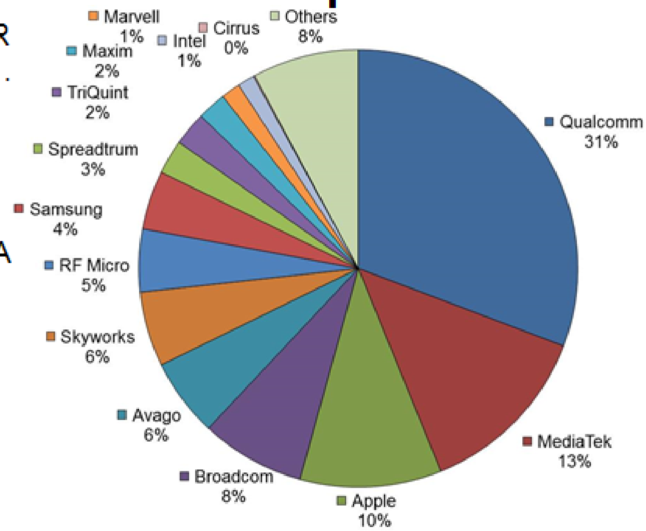

Note: Market share pie excludes memory sold into handsets.

Baseband & AP	***RF & Power Amp***	***Handset Memory***
Qualcomm	Qualcomm	Samsung
Mediatek	Broadcom	Hynix
Intel	Skyworks	Toshiba
Spreadtrum/RDA	MediaTek	Micron
Samsung	Intel	SanDisk
Marvell	Qorvo	

Source: IDC, Gartner, Oppenheimer & Co.

Wireless Handset Semiconductor Forecast: The wireless semiconductor (including memory) market totaled $80.9 billion in 2016, representing 8% growth over 2015. Memory has increased at a 20%+ CAGR over the past five years and is on pace to grow well above that rate in 2017 as favorable pricing conditions have benefitted the market. We expect growth to moderate after a banner 2017 as supply/demand comes more into balance and pricing corrects. With a richer mix of 3G/4G handsets, the move to multi-mode multi-band devices and rising connectivity demands across the spectrum of wireless devices, we believe wireless semiconductor revenue as a whole will rise in line with total handset units, offsetting pricing declines and increased integration. However, there are multiple areas affecting this revenue growth:

- *Phone generation:* 4G phones need more advanced (and more expensive) baseband and RF ICs.
- *Multi-mode capability:* multi-mode phones sometimes need more expensive basebands and additional RF content.
- *Multimedia feature set:* multimedia phones can require applications processors, co-processors, camera modules and additional memory; advanced screens also need more LEDs and battery management solutions.
- *Connectivity:* handsets incorporating Bluetooth, WLAN, GPS, FM and NFC require additional ICs though much of this market is integrating to single chip connectivity solutions.
- *Memory content:* more volatile memory is needed to run advanced software and for non-volatile memory for data storage.

A quick glance points to an increasing semiconductor bill-of-materials. Although handsets remain a consumer-driven market that is extremely sensitive to price, the proliferation of smartphones has created an interesting pricing dynamic. As smartphones increasingly double as mobile computing devices, consumers have shown an increasing willingness to pay for mid-high end devices. And carriers have demonstrated a willingness to subsidize the price of a smartphone so they can collect incremental data usage fees. As the smartphone market continues to penetrate the handset market as a whole, a richer mix of semiconductor content can drive an increasing BOM. However, pricing pressure, particularly in low end emerging market feature phones, lower carrier subsidies over time and increasingly integrated solutions likely will limit the total growth of semiconductor content. Some areas, such as RF content, analog and power management, and application processors will benefit disproportionately to others, such as discrete basebands, and connectivity.

Looking ahead, we would expect the total wireless semiconductor market to grow at a 5% CAGR through 2021; above our assumed roughly flat growth rate in handset unit shipments (primarily a function of migrating to LTE enabled smartphones), but in-line with our ~3% smartphone unit growth forecast. We note that this growth in semiconductor content excludes tablets, which are included in computing but benefit most of the wireless semiconductor vendors rather than traditional PC vendors. In total, we expect handset semiconductor sales to rise from $80.9 billion in 2016 to $104.7 billion in 2021.

In terms of device categories, in 2016 we estimate that discrete basebands represented 6% of IC content, with application processors and integrated BB/AP another 19%. Memory was 33% of revenue. Imaging ICs were another 8%, followed by analog and power management at 11%. On the radio side, RF and power amps represented 16% of total. Connectivity solutions, such as Bluetooth, represented just 7% of IC market revenue in 2016.

Looking ahead, we expect to see some shift in device content, driven by new multimedia and connectivity features and also by integration. We expect both integrated and discrete baseband and application processors to grow at a 3% CAGR through 2021.

Moving to other parts of the baseband sub-system, we expect memory to grow at a 10% CAGR. Analog and power management ICs should grow 3% as some of this functionality is integrated into single-chip SoCs, usually classified as basebands.

On the radio side, we forecast a demand boost from rising band counts in 4G handsets should be offset by initial single-chip cell phone designs that include the RF. 3% CAGR over through 2021 driven by 5% RF IC growth, 1% switching growth, and an offsetting 3% decline in power amplifiers.

Connectivity ICs should decline at a 2% compound rate, arriving at 5% of the BoM 2021.

The exhibits below display handset semiconductor revenue forecast by device type.

Handset Semiconductor Forecast

Baseband (includes memory) **Radio** **Connectivity**

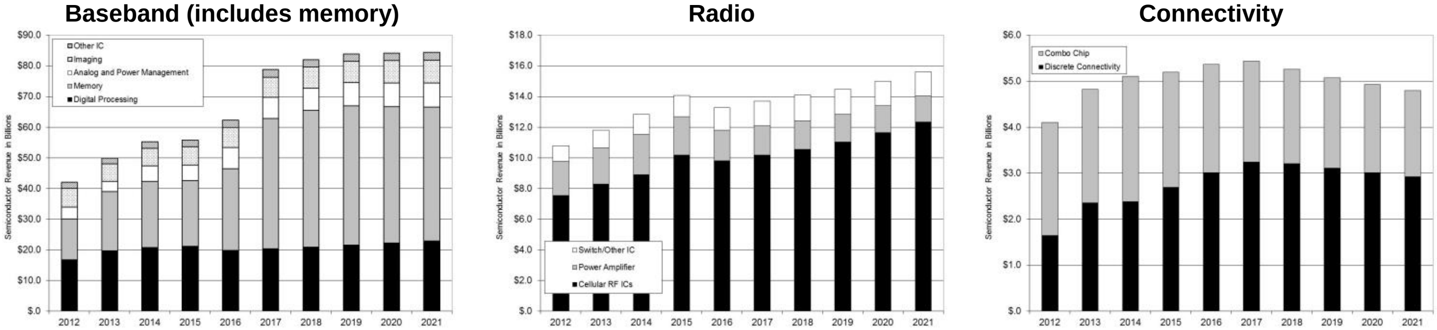

Source: IDC, Oppenheimer & Co.

Wireless Handset Semiconductor Competitors: Semiconductor suppliers in the handset market have undergone substantial changes over the past five years given the rise of Apple, Samsung and low cost providers coupled with the subsequent fall of Nokia, BlackBerry, HTC and LG. Qualcomm has leveraged substantial leadership in LTE and a broad integration strategy into relatively broad share gains while others, such as Texas Instruments, have exited the business entirely. As OEM engineering budgets increasingly fall under pressure with a heightened focus on cost, semiconductor vendors are being looked to for full "systems-engineering" expertise.

Qualcomm was the dominant supplier of handset ICs in 2016, with 31% market share, up from 24% share in 2010. When we published this report in the fall of 2011, Texas Instruments was second in market share with 11% share and has since proceeded to exit the business. MediaTek and Apple were the only other two suppliers with double digit share in 2016 (13% and 10%, respectively).

There are several vendors with 5% share or less, including Broadcom, Apple (its own application processors), Samsung, ST-Ericsson, and Skyworks. Note that this long list still excludes memory suppliers (discussed below).

Looking at the market by subsystem, the discrete baseband market is led by Qualcomm, which held 52% share in discrete baseband components (excluding memory) and 52% of Integrated App Processor/Baseband in 2016, owing to its advantages in CDMA and LTE. MediaTek was No. 2 in the baseband market with 12% share of the discrete segment and in the integrated segment with 27% share. Intel also re-emerged as a competitor in 2016 in the discrete baseband market with a high-profile design win at Apple. Apple continues to lead the discrete baseband processor market with the iPhone the largest non-integrated device.

Note that the relationships between baseband IC vendors and their OEM customers vary greatly. Suppliers like Qualcomm (in some cases) sell packaged chipsets that include software, protocol stacks, and reference designs. Other IC suppliers partner more closely with their customers, customizing their solution to the needs of the individual customers. Note that this latter set of companies also pursues full system solutions with reference designs to serve other OEMs and ODMs outside of their core customer set.

The exhibits below display handset semiconductor revenue forecast by device type.

2016 Baseband and Application Processor Market Share

2016 Discrete Baseband Market Share

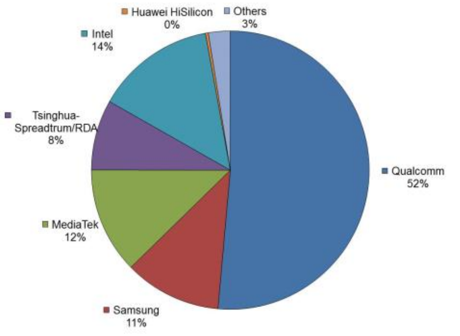

2016 Discrete Application Processor Market Share

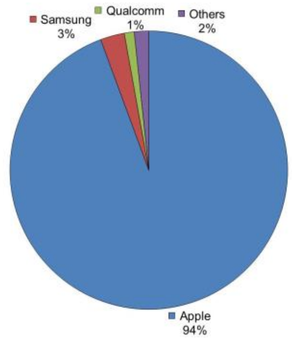

2016 Integrated Baseband/Apps Processor Market Share

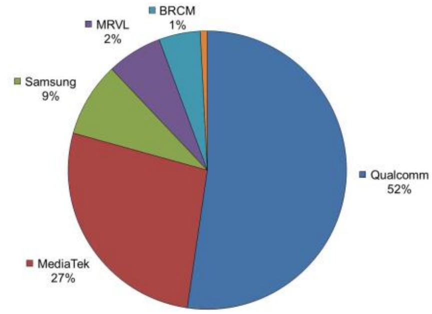

Source: IDC, Oppenheimer & Co.

On the RF and power amp side, the market is somewhat more fragmented. Skyworks, which has extended its PA market share leadership position in recent years, was No. 1 with 27% share. Following Skyworks was Qorvo (combination of RFMD and TQNT), as No. 2 with 23% share.

Qualcomm was the No. 3 supplier overall in 2016 with 21% share, due to its 50%-plus share of the RF transceiver market, where it supplies RF components to go with its basebands. Qualcomm in early 2013 announced plans to enter the PA market with an integrated RF solution, RF360, but to date, has yet to sell any significant volume.

Rounding out the top vendors are Broadcom and MediaTek with 11% and 6% market share, respectively. The exhibits below display RF transceiver and power amp market shares individually.

2016 Handset RF & PA Market Share

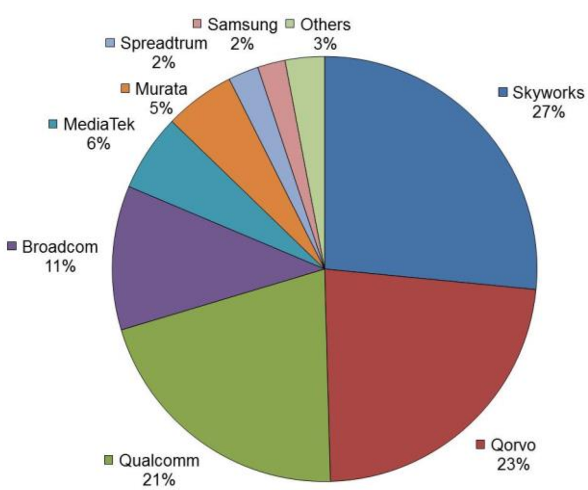

2016 RF Transceiver Module Market Share

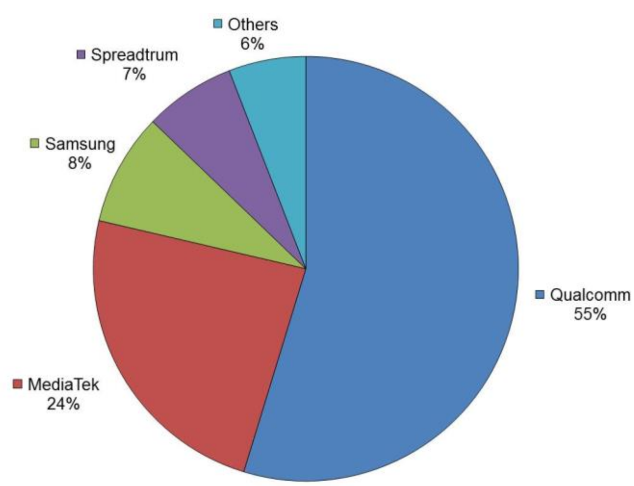

Source: IDC, Oppenheimer & Co.

Also participating in the handset semiconductor market are the vendors of flash memory and DRAM. On the flash side, the largest makers of flash for handsets are Toshiba, Samsung, Hynix and SanDisk. Suppliers of NAND flash are becoming more prevalent in handsets with the growth of high-end smartphones. Samsung has been a driving factor in consolidating market share of both the flash and DRAM markets.

The handset DRAM market is slightly more consolidated with Samsung accounting for 57% market share. The other top suppliers include Hynix and Micron. Note that this includes vendors of multi-chip package memories that include flash and DRAM (many of these vendors source some of their DRAM externally and resell it).

The exhibits below display flash and DRAM market shares for wireless handsets individually.

2016 Handset Memory Market Share

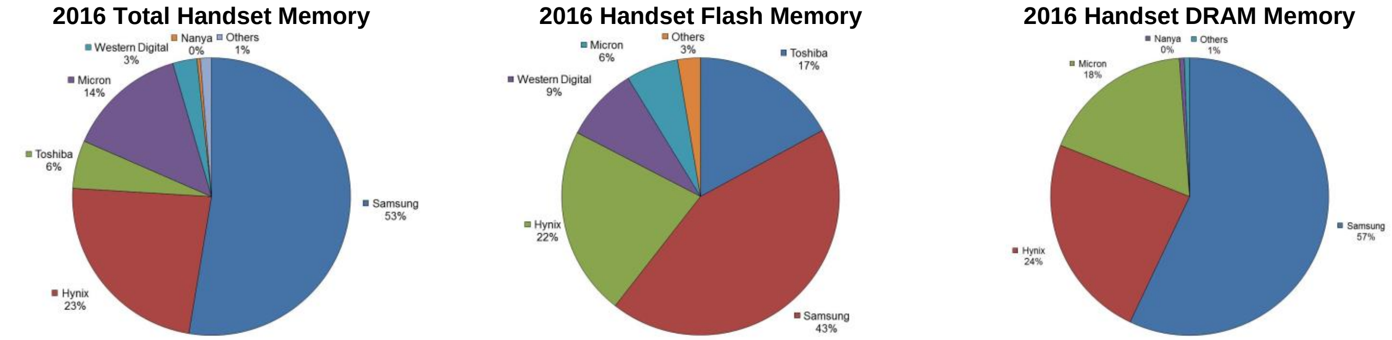

Source: IDC, Oppenheimer & Co.

Wireless Handsets

Wireless Handset Semiconductor Market Trends

- The high end of the smartphone market is maturing and is poised for slower growth going forward:
  - This market is clearly focused on the global transition to LTE
- The feature-phone-to-smartphone upgrade cycle in emerging markets remains in the early-middle innings
- Feature rich handsets require additional application processing:
  - Integrated multi-mode baseband/application processors
- Integrated apps processors, memory and RF front ends are the fastest growing semiconductor segments within wireless handsets.
- Increased complexity driving RF component trends:
  - The move towards multi-mode, multi-band 4G devices
  - RF modules integrating PA, switch, power, and filter etc.
  - Envelope tracking improves performance and efficiency
- Features such as NFC and fingerprint, and facial recognition.
- ODMs remain an integral part of the supply chain with growth in white box market.
- More advanced handsets require more memory and a more efficient battery.

- **The high end of the smartphone market is maturing, is poised for slower growth going forward and is clearly focused on the global transition to LTE.** As the smartphone market grows and has now reached near-saturation in developed markets, consumers have begun to show reduced demand for new, high priced handsets that lack an abundance of incremental features over the prior generation handset. Carriers have shifted from a subsidy model to a payment model, which has pressured the low- to mid-end of the market. As the "flagship" model growth slows and countries (like China) around the world are in the process of rolling out LTE networks, semiconductor vendors must focus on the LTE market opportunity as the next leg of growth, until the next wave of 5G becomes a larger portion of the market in several years.

- **The feature-phone-to-smartphone upgrade cycle in emerging markets.** Always cost-conscious emerging market consumers are upgrading to low-cost smartphones as a means of cheaply acquiring computing power and accessing the Internet. Though there is 4-5x the semiconductor content available in a smartphone than in a traditional feature phone, the effects of this market dynamic are twofold: slowing PC unit growth as a means of primary computing device and pressuring would-be total handset BoM dollars with explosive growth of entry level smartphones. Semiconductor vendors are adapting with full "system level" turnkey solutions to capture a larger BoM share while targeting fast design cycles in the fragmented market.

- **Feature rich handsets require additional application processing and the growth of integrated multi-mode baseband and application processors.** Tied in with the move toward high-performance application processors and basebands supporting multiple cellular modes (3G/4G etc.) are fully integrated chipsets. These solutions have been and will continue to be effective in low-mid-range phones but with rising importance on the LTE market that will likely be the largest cellular standard in history. Though growth here is notably slowing, we would expect discrete apps/baseband processors to remain the choice at the high end of the market where differentiation and performance remain paramount.

- **Connectivity ICs (combo chips) are expected to be flat through 2019.** The number of smartphone units surpassed the number of feature phone units for the first time in 2013. WiFi, Bluetooth, GPS, NFC, and FM, and their integrated counterparts "combo chips" have become embedded in the consumer smartphones. Connectivity is no longer an incremental feature, but an embedded technology, thus ubiquitous and fully penetrated. We model for discrete connectivity solutions to perpetually decline as features such as standalone WiFi are integrated into combo chips.

- **Increased complexity driving RF component trends: the move toward multi-mode, multi-band 3G/4G devices, RF modules integrating PA, switch, power, filter etc., and tunable RF and fully integrated solutions?** As the market moves toward a richer mix of 3.5G/LTE/LTE-A connected devices, the RF front end of the phone is becoming increasingly complex to support a rising number of cellular frequency bands. We believe the RF Front End is one of the fastest growing segments of the wireless market. We see three primary trends here. 1) The move toward multi-mode, multi-band devices is driving more content (PAs, filters, switches, antennas) to support higher band counts. Carrier aggregation to support next generation LTE-A further accelerates necessary RF content. 2) The OEM "system level" approach has bifurcated the market into those suppliers who can provide fully comprehensive and customizable RF modules, and those who cannot. And, 3) because of these aforementioned trends, tunable components are being developed to reduce space and increase efficiency in the front end. Though we do not forecast a quantifiable market will emerge over the next three years, multiple component suppliers are developing monolithic, fully integrated, RF components.
- **New features such as NFC and fingerprint recognition.** Near Field Communications, or NFC, adoption has accelerated, after being driven to prominence by Apple's ApplePay. Integration into connectivity combo chips has also increased prominence in the market. Fingerprint recognition technology has been developed to offset consumer fears of piracy and theft of mobile devices, allowing smartphone users to lock and unlock their device using fingerprint recognition. This market has neared maturity in high-end devices, noted by significant ASP pressures, and is now being replaced by facial recognition in the highest end devices. Further features such as gesture and motion control are likely to continue to proliferate from the high end of the market into mainstream smartphones.
- **ODMs remain an integral part of the supply chain with growth in white box market.** ODM suppliers in Taiwan, China, and Korea have become a very important constituency in the handset market. ODMs play a role similar to OEMs in some cases in that they sell either their own branded phones or sell directly to operators. This "white-box," or non-branded, handset market continues to rapidly grow in emerging economies. From a component perspective, ODMs tend to demand complete chipset solutions for a very quick design cycle to accelerate time-to-market. At the same time, ODMs do not require the same level of customization or performance as OEMs. Most major chipset providers have developed turnkey reference platforms to tailor to the explosive growth of this low-cost handset market.
- **More advanced handsets require more memory and a more efficient battery.** High-end smartphones have dramatically increased the demand for memory capacity, and the constant strain of mobile application use and data streaming has stretched battery life thin. As smartphone form factors on the whole remain largely unchanged yet semiconductor content within the phones increases, the demands on memory and battery are to become more efficient with a shrinking physical footprint. Memory and battery suppliers must constantly pack more capacity/life into the same or smaller form factor.

Wireless Infrastructure

2.5G/3G Wireless Infrastructure System

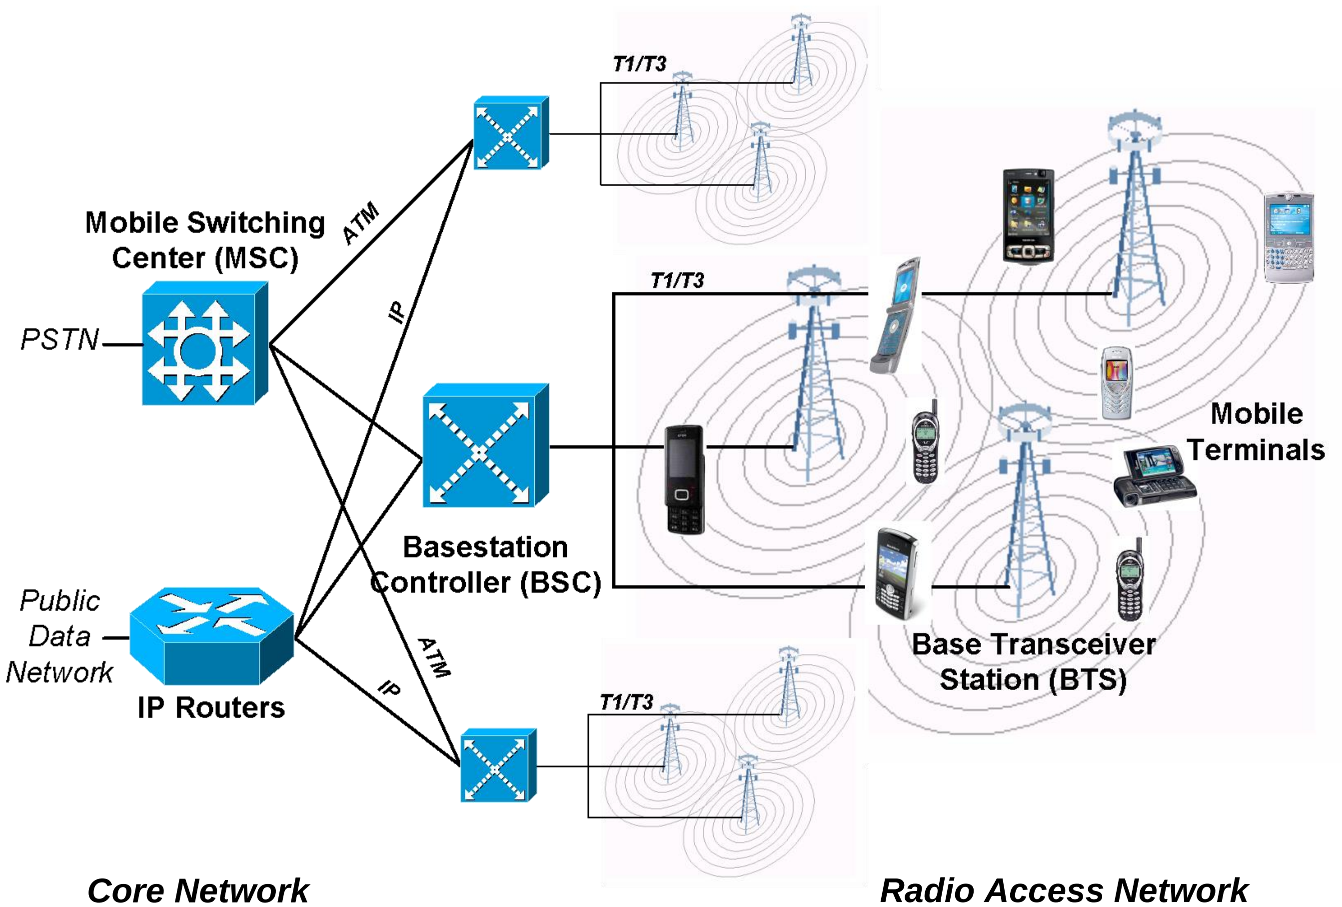

Source: IDC, Computer Desktop Encyclopedia, Oppenheimer & Co.

Wireless voice and data service is enabled by dedicated wireless networks that act as extensions of the public wireline telecom/datacom network. Wireless carriers often overlap their coverage areas to compete with one another; they can also offer service to each other's subscribers. Carriers have acquired wireless spectrum to frame their networks.

The public wireless network is divided into two segments: the radio access network and the core network. The radio access network consists of base transceiver stations (BTSs), often referred to as "cell towers" or just "basestations." Each BTS has a defined range and sends and receives digitally encoded voice and data traffic to and from mobile terminals in its coverage area (sometimes called a cell, as in "cellular"). When a connected terminal passes out of the basestation's coverage area, the call is passed to a nearby BTS (coverage areas generally overlap), allowing for an uninterrupted signal as users move within the carrier's network.

Base transceiver stations are managed by a network of basestation controllers (BSCs). BSCs sit between the radio access network and the core network. They pass traffic back and forth between BTSs and also forward calls to and from the core network. The core network manages the interface between the greater wireless network and the public wired telecom network. Nodes in the core network include mobile switching centers that handle voice traffic (interfacing with the PSTN) and IP routers that handle data traffic (interfacing with the carrier's data network and the Internet).

Wireless Infrastructure

Wireless Infrastructure Revenue Forecast

•We forecast the wireless infrastructure market to moderately increase at a 1% CAGR, from $39.7 billion in 2016 to $41.2 billion in 2019.

•Worldwide LTE buildouts and UMTS (W-CDMA, 3G) spending to satisfy data demand should offset declines in both CDMA and GSM/GRPS.

•LTE/LTE-A are likely to dominate the future of wireless infrastructure builds. We expect LTE spending to rise from 50% of spending in 2016 to 60% in 2021; 2G/3G should steadily decline and represent ~5% by 2021.

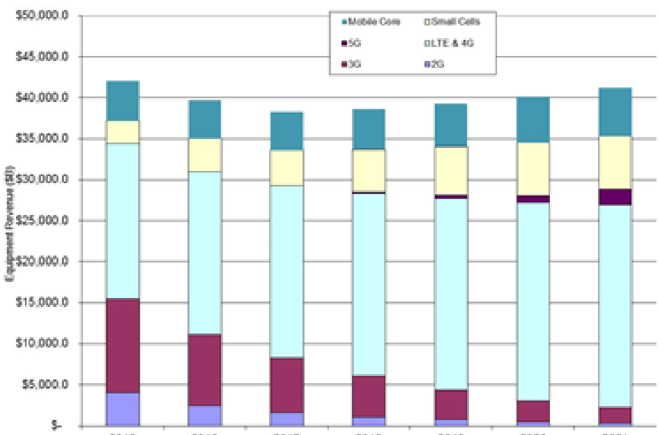

Source: Gartner, IDC, Oppenheimer & Co.

Wireless Infrastructure Suppliers

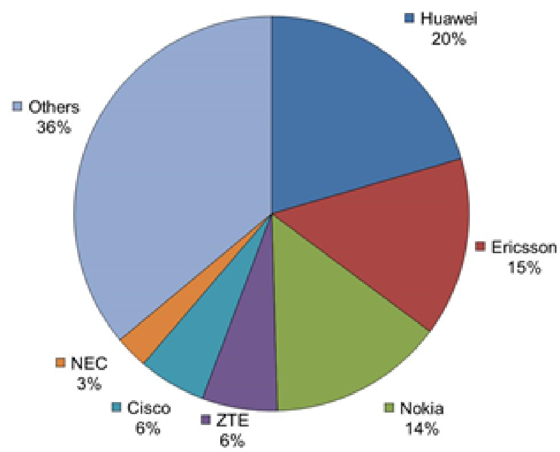

Forecast and market share is on a revenue basis and includes all wireless infrastructure, not just basestations.
Source: Gartner, Oppenheimer & Co.

Wireless Infrastructure Equipment Forecast: Infrastructure spending grew steadily from 2008 through 2011 before beginning to moderate. The sharp increase occurred on two fronts: carriers in emerging markets raced to build out coverage areas, and carriers in North America were rapidly spending on LTE networks. Since this surge, the replacement cycle has begun, with carriers in both developed and emerging markets focused on expanding LTE coverage given its significant cost advantage over traditional 2G and 3G networks.

Looking ahead, we expect the market to grow modestly as network coverage has reached full penetration and growth is likely to come exclusively from emerging markets, until 5G begins to move the needle around 20200. LTE network deployments largely offset declines across 2G, 3G and WiMax. In total, we expect the market to grow modestly, from $39.7 billion in 2016 to $41.2 billion in 2021. The LTE build is well under way, and we expect momentum to continue to accelerate as the pace of adoption increases in emerging markets.

Wireless Infrastructure Suppliers: Huawei was the top supplier of wireless infrastructure in 2016 with 20% revenue share, continuing to take advantage of strong Chinese market growth and a global position (outside of the US). Ericsson was No.2 with 15% share, maintaining a leadership position on LTE infrastructure. Nokia Siemens and ZTE held 14% and 6% share, respectively. Rounding out the top five was Cisco.

Wireless Infrastructure

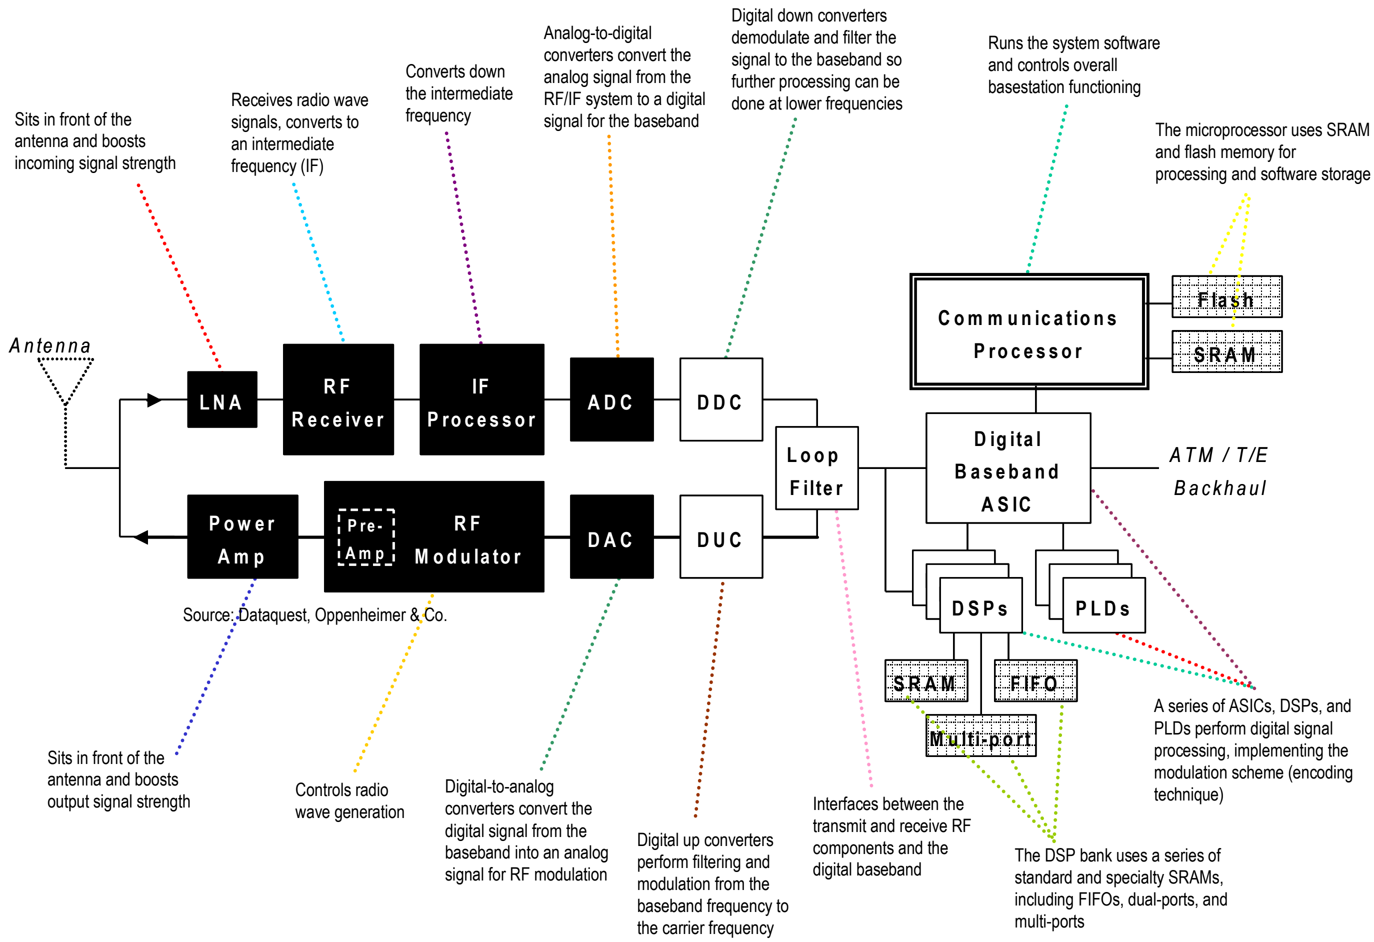

Source: Dataquest, Oppenheimer & Co.

Note: Basestation systems usually separate functions such as RF transmission and reception, power amplification, baseband processing, system control and backhaul transport on separate cards within the basestation. The diagram above depicts a simplified system that would contain all of these functions on a single card—we do not normally find systems like the one depicted above in the field.

Also note that unlike our other block diagrams, most basestation systems are actually less integrated than depicted in our block diagram, often breaking out what is represented in our diagram as a single device into a block of separate devices.

Finally, note that we do not display the ICs that perform backhaul and the interface to the wireline network. Though these systems are a significant portion of wireless equipment spending the devices used in these equipment types are actually more similar to wireline ICs and are therefore included in our discussion of communications infrastructure. This is similar to what we did in the broadband section, including broadband-specific ICs but leaving the network interface ICs out of the discussion at that point.

The semiconductors used in wireless infrastructure systems can vary by protocol and by the OEM's architectural design, but they are largely similar across platforms in terms of their function. Voice and data traffic is transmitted and received to and from individual handsets by the antenna and processed by the RF block, which includes power amp and RF modulation components on the transmit side and LNA, RF receiver and IF processing components on the receive side. A digital baseband subsystem performs protocol processing, translating the voice signals into encoded digital signals for transmission—the baseband subsystem will usually include a series of ASICs, DSPs, PLDs, and memories. Wireline ICs implement the backhaul, ferrying traffic calls back and forth between the PSTN (public switched telephone network) and the basestation, while control plane processors control the overall system.

Baseband and System Control

- **Digital Baseband ASIC**: The baseband processor performs digital processing on the signal, implementing the wireless modulation scheme (the encoding technique, e.g., GSM, CDMA, GPRS, W-CDMA, etc.), encoding and extracting voice signals in and out of the transmission protocol.
- **Communications Processor:** A communications processor manages and controls basestation transceiver system functions. Though only one is shown in our block diagram, each subsystem in a basestation could have its own processor or processors, and a separate control card may be used within the system as well to manage the overall system.
- **DSPs:** Given the number and complexity of high-speed signals processed in a basestation system, discrete DSPs are used in conjunction with the baseband ASIC.
- **PLDs:** PLDs are often used within the baseband system, as the multitude of wireless standards makes late-stage customization critical in base station design.
- **FIFO, Dual-Port, and Multi-Port SRAM:** These specialty SRAMs are used for queuing of signals and as glue logic between the DSPs.
- **SRAM:** SRAM is used by the DSP banks and system control processor.
- **Flash:** Flash memory is used by the system processor.

Transceiver and Power Amp Subsystems—Transmit Side

- **Loop Filter:** The loop filter interfaces between the transmit and receive RF components and the digital baseband.
- **Digital Up Converter (DUC)** (transmit side): A digital converter that performs filter and modulation functions from the baseband frequency up to the carrier frequency.
- **Digital-to-Analog Converter (DAC)** (transmit side): A high-performance digital-to-analog converter converts the digital signal from the baseband system to an analog signal for the RF subsystem.
- **RF Modulator** (transmit side): The RF modulator performs radio wave generation on the transmit side, translating between low-frequency analog signals used in the basestation to high-frequency RF signals used in radio communications. It usually includes a preamp to amplify the signal before sending it to the power amp.
- **Power Amp** (transmit side): One of the most important functions in basestation platforms, power amp ICs are high-power specialized ICs that sit in front of the antenna and boosts signal strength on the way out.

Transceiver Subsystem—Receive Side

- **LNA (Low Noise Amplifier):** The LNA is a power amp that receives the signal from the antenna and boosts it for processing by the system.
- **RF Receiver**: The RF receiver down-converts the radio signal for processing by the system, translating between the high-frequency RF signals used in radio communications and low-frequency analog signals used in the basestation, usually by an IF transceiver.
- **IF Processor**: The IF processor performs all quadrature modulation of “I” and “Q” baseband signals, converting the radio signal down from the intermediate frequency (IF) to an analog signal for use by the system.
- **Analog-to-Digital Converter (ADC)**: A high-performance analog-to-digital converter converts the analog signal from the RF subsystem to a digital signal for the baseband.
- **Digital Down Converter (DDC)**: A digital converter that performs filtering and demodulation so that baseband functions can be performed at lower frequencies.
- **Loop Filter** (also listed on the transmit side): The loop filter interfaces between the transmit and receive RF components and the digital baseband.

LTE / LTE-A

Long Term Evolution (LTE) has taken over as the dominant 4G standard. Wireless standards are constantly evolving to keep pace with explosive data growth. LTE is today's standard, but carrier aggregation enables LTE-Advanced, the evolution of LTE into the first true 4G mobile technology.

Initial commercial implementations of LTE-A began in late 2012 in Russia. LTE-A is fully compatible with LTE (Release 8) technical specifications and must meet or exceed 4G requirements on peak data rates, bandwidth, peak/average spectrum efficiency, delay performance and mobility.

From a network operator perspective, LTE-A is a seamless and cheap upgrade to higher data rates. For the consumer, bandwidth on LTE-A networks is 5x the LTE network and downlink speeds can reach 1 Gbps, 3x the downlink speed on LTE.

While LTE and LTE-Advanced deployments are still underway, carriers and vendors are already working on 5G technology, targeting commercialization by 2020.

LTE-Advanced and Carrier Aggregation. The data and spectrum constraint placed on today's infrastructure present new types of challenges for mobile operators. Global LTE network deployments are ongoing, while LTE-Advanced represents the evolution from today's LTE infrastructure into a true 4G Mobile Communications technology. Carrier Aggregation is the key technical enabler of LTE-A and substantially higher data rates, efficient utilization of fragmented spectrum and support of heterogeneous networks. However, each of these changes to the network presents a new challenge for mobile semiconductor vendors. In our view, the biggest beneficiaries of more complex network standards are the wireless RF vendors.

There were three primary motivations for Carrier Aggregation (CA) in Release-10—support of higher data rates, efficient utilization of paired spectrum and support of heterogeneous network deployments. Insufficient peak data rates and spectrum capacity limitations are two of the biggest challenges today facing network operators—CA is the most important feature for the enablement of LTE-A bandwidth specifications and for the maximization of fragmented spectrum capacity. In CA, multiple component carriers (carriers) may be aggregated together as a means of increasing both peak data rates and throughput. LTE Release-10 allows for this aggregation. Each carrier continues to utilize the Release-8 structure, thus maintaining backwards compatibility to ensure seamless crossover.

CA is highly demanded from a network-operator perspective because it allows for the aggregation of spectrum fragments, an avenue away from the spectrum crunch and more efficient (cheaper) network operation. To the consumer, CA means higher data rates. Because carriers remain backwards-compatible with LTE Release 8/9 standards, the upgrade cycle is relatively seamless and LTE networks can smoothly migrate to LTE-A. Both TD/FD-LTE are supported and the relative capex cost to LTE is cheaper. Because of this backwards compatibility, a majority of technical change occurs in the end device rather than the terminal (base station).

LTE / LTE-A

LTE/LTE-A

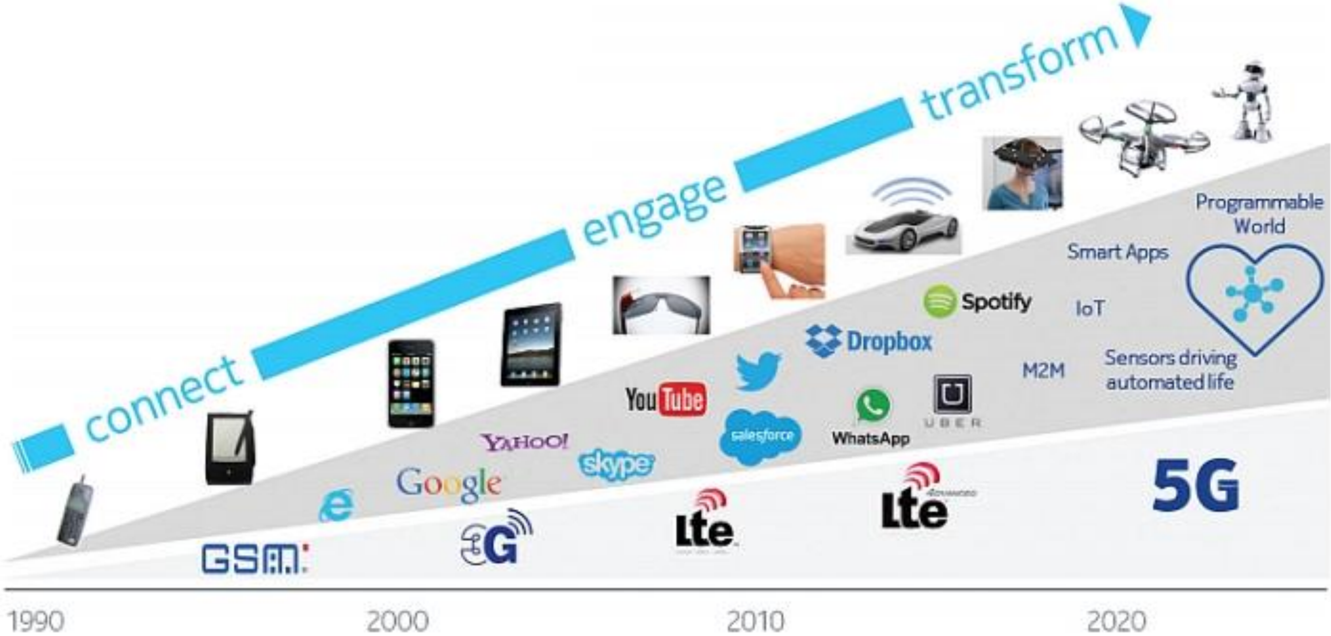

Source: Nokia, Oppenheimer & Co.

Carrier Aggregation

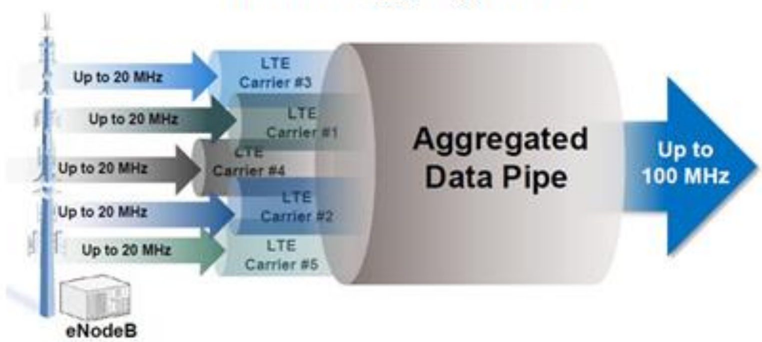

Source: Qualcomm, Oppenheimer & Co.

LTEvolution: OEMs are increasingly seeking semiconductor vendors that can fully customize a handset design at the system level, while carriers want a handset that can maximize network efficiency. We believe this dynamic increasingly puts power in the hands of a few. Front end architectures are today very different. They vary across regions to support different bands and protocols. They vary across carriers. Each of the different front end architectures must support a varying number of frequency bands across geographies and/or carriers, or become an even more complex “world phone” to support north of 40 LTE bands. We aren’t convinced that a single architecture matters more than another—the systems vendor must have the ability to support and design all.

In our view, the biggest wireless beneficiaries of advanced network evolution are the wireless RF vendors to seamlessly enable the more complex wireless network connection. To take the “systems approach” is to have the full suite of leading-edge product offerings that encapsulate the full RF suite across any number of processes (GaAs, SiGe, CMOS, SOI). Based on more band support—and more complex band configurations—demanded by Carrier Aggregation/LTE-A, the emergence of power supply modulation (envelope tracking) and necessity for new technologies like antenna tuning, the RF front end has become more intertwined with, and more reliant on, the baseband and transceiver than ever. Traditional RF vendors stand to benefit from the secular TAM expansion. Baseband and RF transceiver vendors should benefit from integration opportunities going forward.

Digital Set-Top Boxes

Set-top boxes have enjoyed strong growth as the key enablers of premium video services in the home. In developed nations, however, set-top boxes are increasingly being cannibalized by OTT streaming services which are driving “cord cutting.”

In addition to service-based units which include features such as DVR, some set-top boxes are sold at retail directly to the customer. Free-to-air boxes allow consumers to watch terrestrial digital broadcasts—this is common outside the US. Standalone DVR has become a niche, as most service providers now offer these directly to their subscribers.

Set-top boxes have enjoyed strong growth as the key enablers of premium video services in the home. In developed nations, however, set-top boxes are increasingly being cannibalized by OTT streaming services which are driving “cord cutting.” Cable, satellite, and IPTV service providers are all deploying advanced set-top boxes to their subscribers to enable premium features like DVR. In addition to these service-based units, some set-top boxes are sold at retail directly to the customer. Free-to-air, or “digital terrestrial” as it is sometimes called, is common outside the US.

Standalone DVR has become a niche, as most service providers now offer these directly to their subscribers. TiVo and others still sell standalone DVRs, but they tend to serve consumers with specific tastes or needs.

Digital Set-Top Boxes

Set-Top Box Unit Forecast

•We expect set-top boxes to grow at a 3% CAGR, from 288 million units in 2016 to 341 million units in 2021.

•Cable should decline by 2% each year as cord cutting cannibalizes units

•Satellite remains an important market, and likely outpaces industry growth at a 5% CAGR, as the cost of satellite systems has come down, and limed required infrastructure makes it particularly appealing in emerging markets.

•IPTV has become increasingly more important. We forecast units to grow at a 8% CAGR through 2021, reaching 53% of units.

•We forecast Free-to-air to decline 17% annually, as a lack of internet support confines free-to-air set-top box demand to emerging markets.

Key Set-Top Box Suppliers

Satellite	***Cable***	***IPTV***
Technicolor	Arris	Sagem
Echostar	Cisco	Arris
Tongda	Samsung	Cisco
Humax	Technicolor	Samsung

Source: Gartner, Oppenheimer & Co.

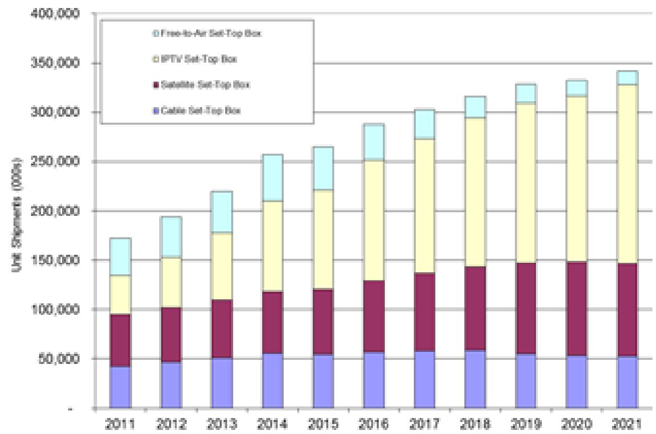

Source: Gartner, Oppenheimer & Co.

Digital Set-Top Box Unit Forecast: The market for digital set-top boxes (STBs) has experienced rapid growth during the past few years, as a number of new technologies pervaded the market, driving strong upgrade cycles across multiple service providers. In total, we estimate the market totaled 288 million units in 2016. Growth is slowing, however, as streaming TV drives cord cutting in developed nations with faster Internet service.

Service providers are migrating customers from analog-based services to digital services, Analog boxes are being wound down replaced with higher-cost and higher-function HD STB's.

In satellite, providers moved to 8PSK for the front end and MPEG-4 encoding for the back end to save bandwidth; this required new set-tops. The satellite providers in the US have conjoined the MPEG-4 and HD migrations in order to allow HD services without adding material bandwidth (satellite launches are very expensive).

In Asia, and US IPTV deployments have accelerated as consumers become focused on obtaining Internet access vs. traditional Pay TV. After a long development period, IPTV has begun to gain momentum in Europe and South American. IPTV remains the long-term solution for telecom providers looking to add video services to drive revenue growth; these services require an entirely new rollout of set-top boxes.

Finally, free-to-air digital set-tops remain a large market internationally, allowing consumers to receive digital broadcasts without a service fee.

Then there is the trend toward digital video recorders (DVRs), which have enjoyed tremendous consumer reception and also make the carrier happy because they can charge additional fees. Though DVRs started as standalone boxes a few years ago, they have been embraced by all of the major cable and satellite carriers and are a part of the IPTV rollout as well. Service-based DVRs now dwarf those sold at retail for either standalone or terrestrial use.

Overall, we estimate the market is growing at a 3% CAGR, from 288 million boxes in 2016 to 341 million in 2021. Cable should decline at a 2% compound rate, driven by HD and DVR upgrade cycles in developed regions that are offset by growth of IPTV-based systems in China and streaming services. According to market researchers Gartner, satellite should enjoy strong growth of 5%, boosted by HD and DVR demand in markets where laying cable or fiber infrastructure is challenging. IPTV should grow at an 8% CAGR through 2021, while free-to-air should decline at a 17% compound rate. The exhibits below display actual and forecasted set-top box unit shipments for each market.

Cable Set-Top Box Unit Shipments

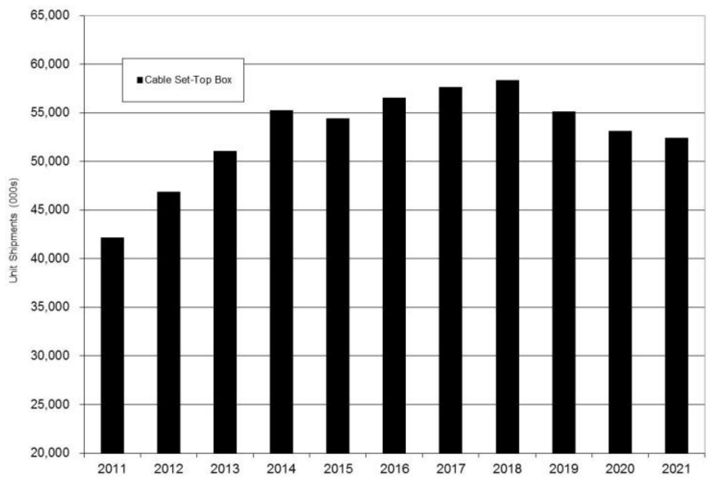

Satellite Set-Top Box Unit Shipments

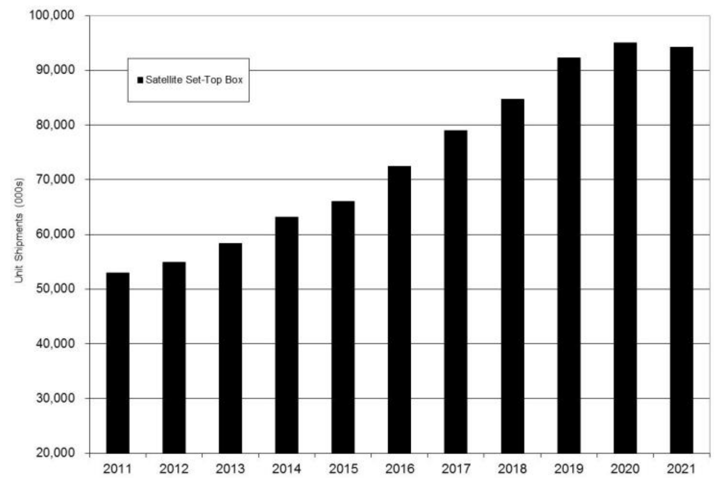

IPTV Set-Top Box Unit Shipments

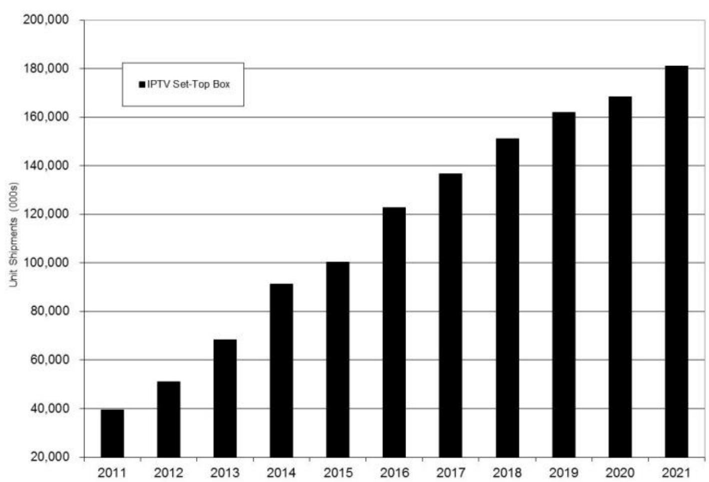

Free-to-Air Set-Top Box Unit Shipments

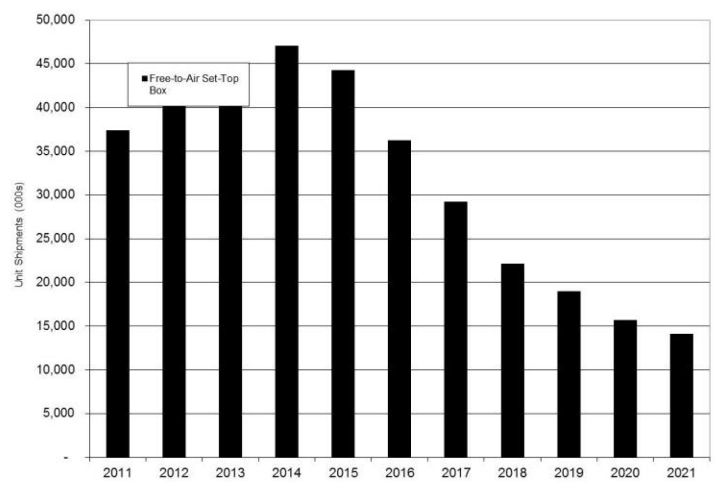

Source: Gartner, IDC, Oppenheimer & Co.

Digital Set-Top Boxes

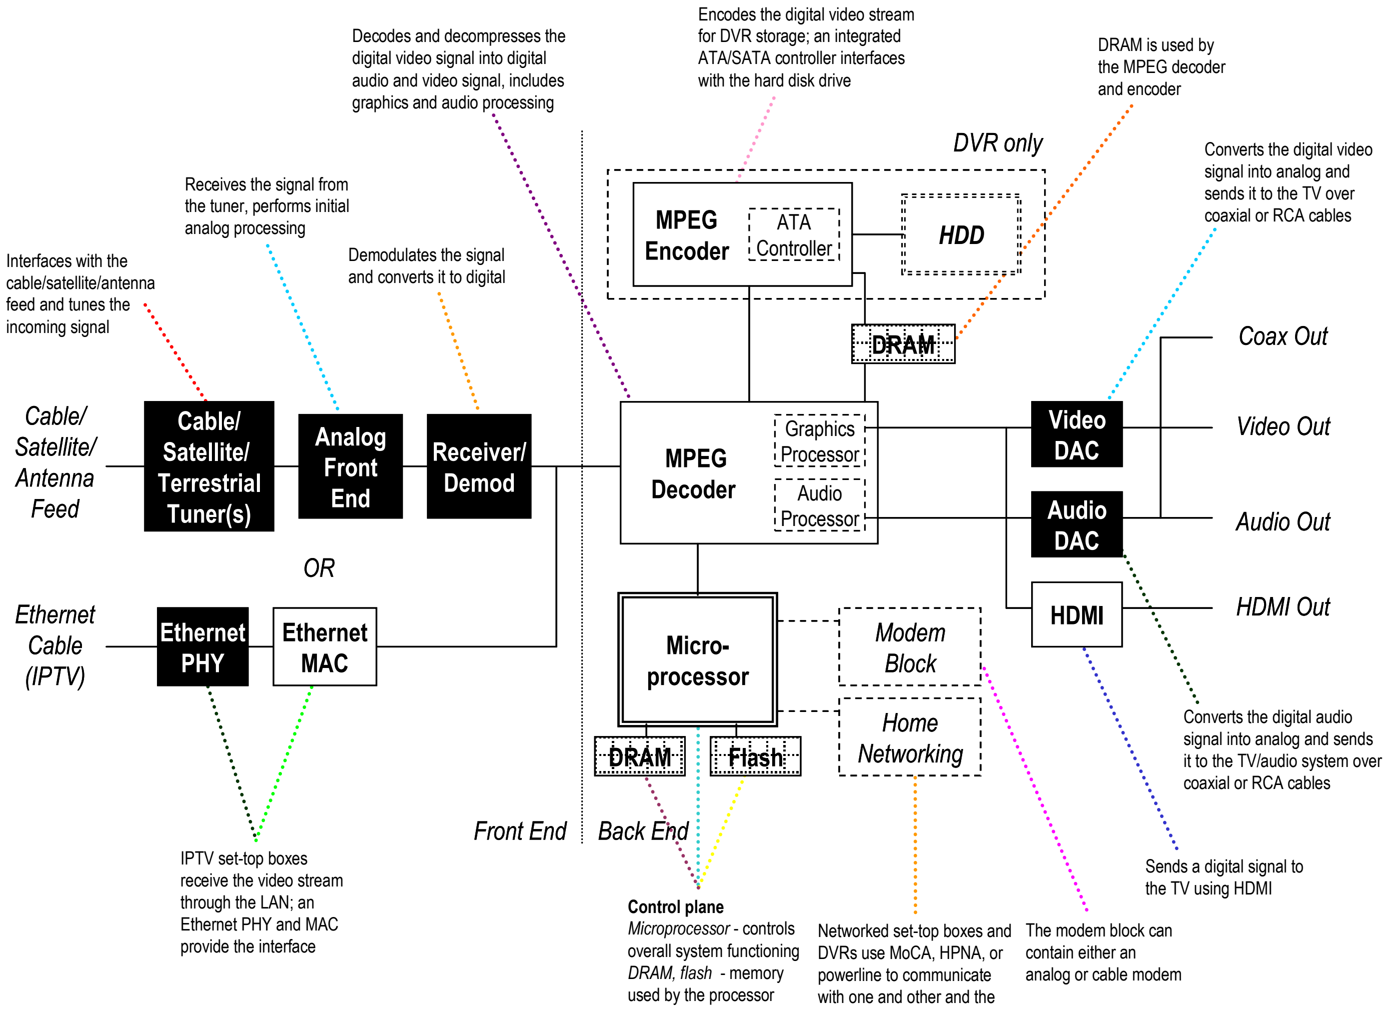

Source: IDC, Oppenheimer & Co.

A set-top box looks very much like a modem with additional audio and video processing functionality. Front end ICs (tuner, AFE, demodulator) receive the signal from the network, and back end ICs (MPEG decoder, graphics and audio processors) decode it. The key difference is that instead of a data stream, the set-top decodes a video signal.

Set-top box semiconductors blend devices from the communications and digital video worlds. A tuner receives and tunes the signal from the coaxial cable, which is processed by an analog front end and demodulator, which converts the analog signal to digital. The signal is then decoded by an MPEG decoder, which can then adjust and add to the stream with integrated graphics and audio processors. The decoded signals are then sent to the television either digitally using HDMI or in an analog format over RCA or component cables (the latter requires a conversion back to analog).

In a DVR, additional ICs are added in to support MPEG encoding as well as hard disk drive interface control. ICs found in the hard disk drive are not discussed here, as they are captured in our discussion of HDDs earlier in the report.

Front End

- **Tuner**: The tuner receives the signal from the coaxial cable, satellite feed, or terrestrial antenna and tunes it to the proper RF frequency. Most set-top boxes have just a single tuner, though some boxes have multiple tuners to support picture-in-picture or other advanced features. DVRs require at least one additional tuner.
- **Analog Front End**: The analog front end receives the signal from the tuner and performs the front-end analog processing.
- **Receiver/Demod**: The receiver chip demodulates the signal and converts it to digital for decode and processing by the set-top box. It is sometimes integrated with the analog front end for an integrated front end ASSP.

- **Ethernet PHY and MAC** (IPTV set-top box only): IPTV set-tops receive a packetized video stream over the LAN as opposed to a direct feed. They therefore do not have tuners, AFEs, or receivers, but add Ethernet ICs to provide the LAN interface. The PHY interfaces with the Ethernet cable; the MAC performs protocol processing.

Back End

- **MPEG Decoder**: The MPEG decoder decodes and decompresses the encoded digital signal into audio and video signals. Most set-top boxes use the MPEG-2 standard, though newer satellite and IPTV boxes are moving to MPEG-4, which offers tighter signal compression.
- **Graphics Processor/Encoder**: The graphics processor performs digital image generation, adding graphic items associated with the interactive guide, on-demand menus, and other services. It also includes an encoder to perform any additional processing on the video signal before it is sent to the TV, such as NTSC, PAL, etc. It is usually integrated into the MPEG decoder.
- **Audio Processor**: The audio processor is a DSP that performs digital audio generation, as well as any additional audio encoding, such as Dolby Digital, DTS, etc., also usually integrated into the MPEG decoder.
- **Microprocessor**: A microprocessor performs control plane functions and runs the set-top box software. It is sometimes embedded in a back-end ASSP, though more advanced boxes will use a discrete processor.
- **Flash/DRAM**: The MPEG decoder and microprocessor use a variety of memories.

Back End Interface

- **Video D/A Converter**: A video digital-to-analog converter converts the video signal to analog format to send it to the TV over coaxial or RCA or component cables.
- **Audio D/A Converter**: An audio digital-to-analog converter converts 16-bit stereo samples into left and right alternating-current audio signals and sends them to the TV over coaxial, optical, or RCA cables.
- **HDMI Interface IC:** More advanced set-top boxes include an HDMI (High Definition Multimedia Interface) interface, which is an all-digital, high-definition standard that supports both video and audio with digital rights management capabilities. An HDMI IC is needed to support this interface.

Modem

- **Analog Modem SLIC and Codec** (not pictured in diagram): Some set-top boxes incorporate the ability to send upstream signals (such as an order for a pay-per-view program) to the service provider over an analog modem. The SLIC drives the signal over the RJ-11 telephone line and the codec performs the digital-to-analog and analog-to-digital conversion for the modem.
- **Single-Chip Cable Modem** (not pictured in diagram): Some newer high-end digital cable set-top boxes incorporate cable modem functionality by adding a single-chip cable modem device and an additional tuner.

Digital Video Recorder (DVR)

- **MPEG encoder**: An MPEG encoder chip converts video signals back into the MPEG format for storage on the hard drive. The encoder can be implemented as a separate device or it can be integrated in with the MPEG decoder.
- **ATA/SATA Controller**: This controller provides the interface to the hard disk drive. It is usually integrated into the MPEG encoder. Some DVRs support an external SATA HDD in addition to the internal HDD.

Home Networking

- **MoCA, HPNA, Powerline Chipset** (not pictured in diagram): Newer set-top boxes contain home networking capability to allow the set-top boxes to communicate with one another and with the home gateway. MoCA uses coax, HPNA uses phone line and coax, and HomePlug, UPA and HDPLC uses powerline. These are discussed in more detail in the emerging technologies section at the back of this report.

Digital Set-Top Boxes

Set-Top Box Semiconductor Forecast

•With strong unit growth and features like HD and DVR boosting the BOM, set-top box semiconductor revenues have increased by almost 60% in the past five years.
•We expect a more modest 3% CAGR through 2021.
•IPTV should lead the way, reaching $4.6 billion in 2019, an 8% CAGR.
•Satellite should grow at 4% and cable should decline 1%, while free-to-air should be pressured by declining ASPs and lower unit growth.

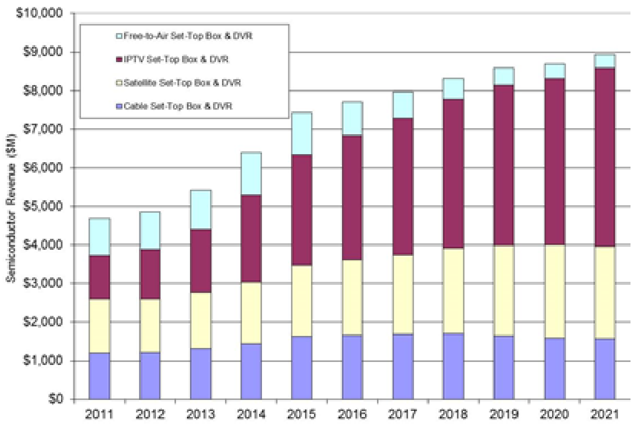

Note: Excludes DVR/DVD-R boxes and place-shifting units without a tuner (e.g. SlingBox).
Source: Gartner, Oppenheimer & Co.

Set-Top Box Semiconductor Competitors

STMicro
Broadcom
Intel
MaxLinear/Entropic
Ali Tech
Sigma Designs
Montage
ViXS
Sony
Samsung

Source: IDC, Oppenheimer & Co.

Digital Set-Top Box Semiconductor Forecast: With strong unit growth and features like HD and DVR boosting the semiconductor bill-of-materials, set-top box IC revenues have grown almost 60% in the past five years. Satellite and IPTV were the primary drivers, with cable contributing as well. In 2016, the market totaled $7.7 billion. Looking ahead, we expect content gains to be partially offset by normal price pressure, and thus forecast revenues to rise at a 3% CAGR to $8.9 billion in 2021.

Digital Set-Top Box Semiconductor Competitors: The set-top box IC market is highly concentrated among the top two vendors: Broadcom and STMicro. Broadcom has been a top provider in cable (dominates Arris), satellite (Echostar, DirecTV, Sky Networks) and IPTV. STMicro has been a dominant player in satellite (DirecTV, Echostar, international) and in free-to-air, and a significant player in cable (Cisco). Other vendors include MaxLinear (through Entropic), Ali Tech, IPTV specialist Sigma Designs, Montage, ViXS, Samsung, and Sony.

Automotive

Automotive Electronics Systems

Airbags

Anti-lock Brakes

Advanced Key Systems

Engine Control Management

Infotainment

Advanced Driver Assistance Systems

Over the last decade, the automotive industry has been marked by a period of technical innovation as impactful as any in its history. Auto makers are increasingly focused on safety, reliability, fuel efficiency, and performance, and are relying on advancement in semiconductor technology to achieve their goals. Looking forward, the automobile will continue its metamorphosis from a “dumb” mechanical device to a “smart” electronic one, driving sustained semiconductor content growth. In 2016, the automotive semiconductor market represented $33.1B, or 10%, of the $340B semiconductor market. We conservatively expect automotive-related semiconductor revenue to grow at a 7% CAGR to $47.3B by 2021.

Key electronics systems included in modern automobiles include:

- Advanced Driver Assistance Systems (ADAS)
- Airbags (drive, passenger and side airbags)
- Antilock braking systems (ABS)
- Auto stereo (AM/FM radio, cassette player, CD player and changer, satellite radio)
- Multi-zone climate control
- Dashboard instrument cluster / Infotainment
- Engine control
- Remote/keyless entry
- Autopilot / Self Driving
- Battery Management Systems
- In-vehicle networking

Automotive

Auto makers are increasingly focused on safety, reliability, fuel efficiency, and performance, and are relying on advancement in semiconductor technology to achieve their goals. Looking forward, the automobile will continue its metamorphosis from a "dumb" mechanical device to a "smart" electronic one, driving sustained semiconductor content growth. In 2016, the automotive semiconductor market represented $33.1B, or ~10%, of the $340B semiconductor market.

Automotive Electrical Systems/Vehicle

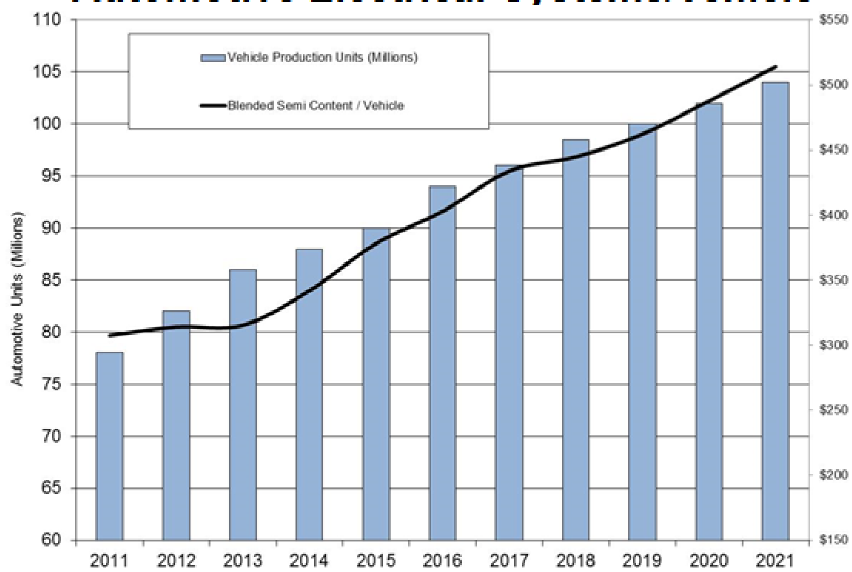

Source: Gartner, Oppenheimer & Co.

As a mature market, automobile units grow at a low single-digit rate annually. Average semiconductor content, however, is expected to outgrow vehicle units at least 2x over the next decade. In 2016, we estimate the average small-sized vehicle had ~$165 of semiconductor content, while mid-sized averaged ~$462 and luxury averaged ~$910. By 2021, we estimate average semiconductor content in small vehicles will be $215; in mid-sized, $610, and in luxury, $1,050.

Demand for greater fuel efficiency, lower emissions, and higher performance are all driving higher levels of semiconductor content within the powertrain. Automotive OEMs are now designing safety systems which utilize a variety of advanced semiconductor and sensor technology to focus on accident prevention, rather than damage mitigation. The proliferation of these systems has been driven largely by two seismic shifts: government regulation/insurance initiatives and the influence of safety features on consumers' purchasing decisions.

Infotainment semiconductor content makes up ~32% of automotive semiconductor content today. We believe OEMs seeking product differentiation will increasingly lean upon suppliers of advanced graphics, and processing systems for their infotainment systems. This is expected to drive significant semiconductor content growth in support of various connectivity standards (Bluetooth, Wi-Fi, NFC, etc.), power conversion/management needs, and displays.

Automotive

Automotive Electronics Dollar Content per Vehicle Forecast

•We expect automotive semiconductor content per vehicle to grow at a 5% CAGR, rising from $403 per vehicle in 2016 to $514 per vehicle in 2021.

•Infotainment semiconductor content made up ~32% of automotive semiconductor content in 2016.

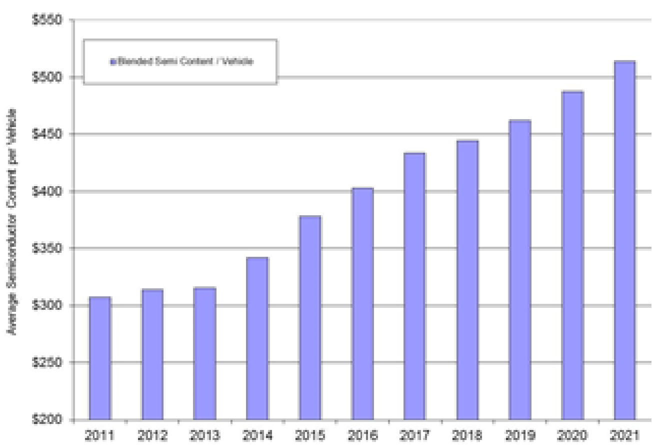

Source: Gartner, Oppenheimer & Co.

Automotive Electronics Suppliers

Americas

Delphi Automotive
Johnson Controls
Garmin
TRW
Visteon
Lear Corporation
Harman
Arvin Meritor
Eaton
Stoneridge
Key Safety Systems
Magellan

Europe

Robert Bosch
Continental
TomTom
Autoliv
Magneti Marelli
Hella
Valeo Electronics
Kostal
Marquardt
Haldex
Behr-Helia
Kiekert Group

Japan & Asia

Denso
Aisin
Mitsubishi Electric
Panasonic
Hitachi Automotive
Pioneer
JVC Kenwood
Calsonic Kansei
Alpine Electronics
Mitac
Fujitsu-Ten
Clarion
Takata
Hyundai Autonet
Alps Electric
Sanyo

Source: Gartner, Oppenheimer & Co.

Automotive Electronics Unit Forecast: In 2016, automotive-related sales represented $33.1B, or ~10%, of the $340B semiconductor market, up 10.0% from $30.1B in 2015. Rapid technological innovation, regulatory demands, and the metamorphosis of the automobile from a mechanical device to an ever more electronic one have given rise to one of the longest-tailed, most stable drivers of semiconductor growth in recent memory. In total, we expect automotive semiconductor content to grow at a 7% CAGR from 2016 to 2021, reaching $47.3B in 2021. On a per vehicle basis, we expect automotive content to grow at a 5% CAGR from 2016 to 2021, to $514 per vehicle.

Automotive Electronics Suppliers: Top suppliers of automotive electronics include Robert Bosch, Continental Automotive, Denso, Delphi Automotive, Johnson Controls, Garmin, TomTom, Aisin, Autoliv, Mitsubishi Electric, TRW, and Visteon. Note that market share in each geography is highly skewed toward regionally located vendors; we display top vendors by region in the slide above.

Automotive

Automotive Semiconductor Forecast

- We expect the automotive semiconductor market to increase at a 7% CAGR, growing from $33.1 billion in 2016 to $47.3 billion in 2021.
- ADAS should see the highest growth, rising at a 17% CAGR. All other segments should grow at a ~7% CAGR with the exception of body/chassis, which should decline at a 2% compound rate.
- Simpler electronic systems still rely heavily on microcontrollers, while more complex systems use MCUs along with ASICs, ASSPs, DSPs, and FPGAs.

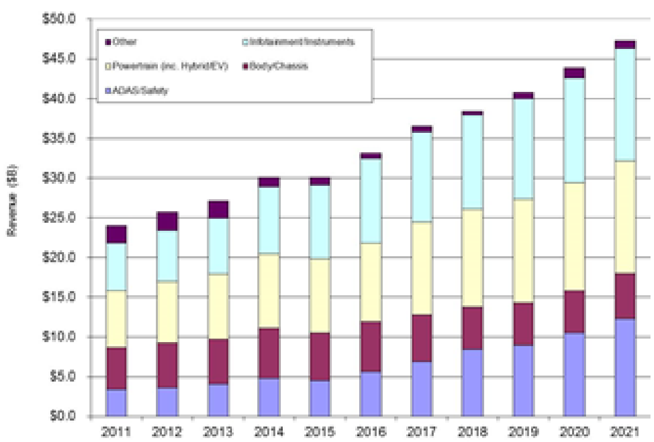

Source: Gartner, Oppenheimer & Co.

Automotive Semiconductor Competitors

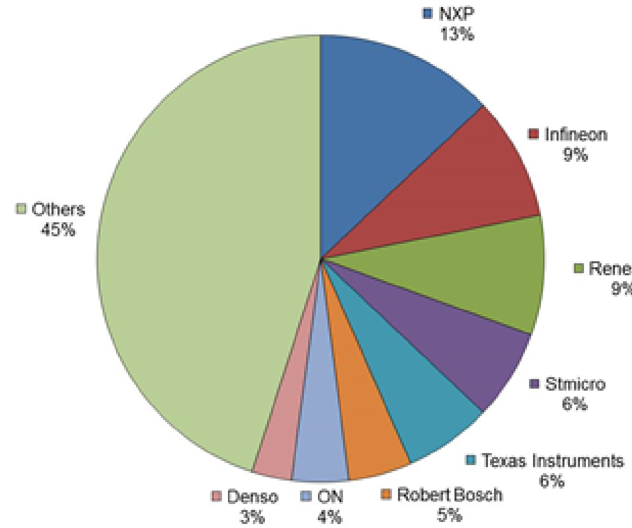

MCUs	ASIC/ASSP	Analog	Discretes
Renesas	STMicro	Sanken	Renesas
NXP	Infineon	ADI	Infineon
TI	NXP	ON Semi	Rohm
STMicro	Robert Bosch	Maxim	ON Semi
Denso	Renesas	TI	NXP
Infineon	TI	Renesas	Denso
Fujitsu	Panasonic	Rohm	Toyota
Microchip		Toshiba	
Toshiba			

Note: Market share includes automotive microcomponents, ASICs/ASSPs, analog ICs, and discretes.
Source: Gartner, Oppenheimer & Co.

Automotive Electronics Semiconductor Forecast: We believe we remain in the early stages of the car's metamorphosis from a dumb mechanical device to a smart electronic one, driving both increasing complexity of each automotive electrical system and a slight mix shift toward systems with greater semiconductor content, such as ADAS and autopilot/self-driving cars.

We estimate that the automotive semiconductor market represented $33.1B, or ~10%, of the $340B semiconductor market in 2016. We forecast the market to increase at a 7% CAGR through 2021 to $47.3B. Our work suggests ADAS/safety applications will see the fastest growth in content over the next five years, growing at a 17% CAGR from $5.6B in 2016 to $12.3B in 2021. Powertrain, infotainment and other are all likely to grow at ~7% CAGR, with the exception of body/chassis, which is likely to decline by a 2% compound rate.

Automotive Electronics Semiconductor Competitors: Automotive semiconductors are supplied primarily by the larger, diversified semiconductor suppliers. In 2016, only NXP, which was formed through the merger between Freescale (7% share in 2014) and NXPI (6% share in 2014), held a market share greater than 10%.

Automotive

Automotive Semiconductor Market Trends

- Given the high-performance and sticky nature of the automotive semi market, we expect this to drive profitable growth and sustained visibility for those leveraged to the segment.
- Demand for greater fuel efficiency, lower emissions, and higher performance are all driving higher levels of semiconductor content within the powertrain.
- We believe the infotainment market has seen growth slow as features have become more pervasive in new vehicles, with growth primarily from incorporating ADAS functions in infotainment systems going forward.
- Automotive OEMs are now designing safety systems which utilize a variety of advanced semiconductor and sensor technology to focus on accident prevention, rather than damage mitigation.
- We believe we remain in the early stages of the car's metamorphosis from a dumb mechanical device to a smart electronic one. This trend likely drives sustained, high profit growth for semiconductor companies for years to come.
- While mass adoption of driverless cars appears unlikely for at least 5+ years, we believe major paradigm shifts in safety, fuel efficiency, and entertainment foreshadow a long and sustainable growth opportunity for automotive semiconductor suppliers.

Emerging Technologies

Key Emerging Technologies:

Artificial Intelligence

100G

Antenna Tuning / Control

Picocells and Femtocells

Wearables / Smart Watches

LoRa – Industrial IoT Connectivity

This section deals with emerging technologies, those where the market is just beginning to come together and volume deployments are about to begin or are in the very early stages of ramping. It is likely that many of these technologies will move from this section closer to the front of the document in future updates; in fact, many of these are already mentioned peripherally in earlier sections.

For each section, we attempt to give just a brief introduction to the technology, its targeted application, and the major semiconductor players to watch.

Artificial Intelligence

Artificial Intelligence and Machine Learning

Artificial intelligence (AI), at its base, is defined as intelligence exhibited by machines. This intelligence is defined by a machine's ability to take inputs and, through an array of instructions and processes, generated desired outputs.

Machine learning entails having a machine generate code and algorithms through learning from previous inputs and outcomes to analyze data - while simultaneously correcting algorithmic errors - and make predictions based on what it has learned.

- If imagined as a classroom, machine learning would be similar to giving students a textbook and numerous practice tests before asking them to handle real world applications.

Source: nVidia

Source: Intel

Artificial intelligence (AI), at its core, is defined as intelligence exhibited by machines. This intelligence is defined by a machine's ability to take inputs and, through an array of instructions and processes, to generate desired outputs. As current trends and advancements show, machine learning is quickly becoming capable of tackling more difficult and complex situations, marking the true potential of AI and its applicability to numerous roles in everyday life.

Recently, artificial intelligence has been aided by two fundamental shifts:

1) A shift to viewing variables and solutions as more flexible, and less binary, in nature
2) A shift to delegating more artificial intelligence development responsibilities to machines, facilitated through **machine learning**, as opposed to human intervention**.**

Machine learning entails having a machine generate code and algorithms through learning from previous inputs and outcomes to analyze data—while simultaneously correcting algorithmic errors—and make predictions based on what it has learned.
If imagined as a classroom, machine learning would be similar to giving students a textbook and numerous practice tests before asking them to handle real-world applications.

We see several trends driving the recent acceleration in AI adoption and development:

- **Cost of computing power has decreased:** Semiconductors will likely continue the trend of becoming more power- and cost-efficient.

- **GPU application to AI:** Parallel processing facilitated by GPUs allows many simple tasks to be complete quickly and simultaneously.

- **Availability and pervasiveness of data collection and storage:** In the age of big data and interconnected networks, the ever-increasing availability of data has facilitated a foundation for AI development.

- **Move to hyperscale and public cloud centralizes computing power data storage, and software developers to utilize the data:** Firms will likely increase their reliance on cloud computing for increased performance, IT agility and other associated services.

- **Enterprise spending on AI:** As costs continue to fall and firms realize the wide applicability of AI and potential ROI, subsequent AI investment will likely follow.

- **Funding and investment in AI:** Public and private companies alike will look to fund companies generating AI-related products or services, thus spurring future innovation.

Artificial Intelligence

Neural Networks

•Machine learning leads to the generation of **neural networks** that provide a procedure to analyze presented data. Deep learning utilizes various levels of interconnected neural networks.

- **Neural networks** function similar to the system of neurons within a brain – numerous **nodes**, the machine learning equivalent of neurons, accept, process and rely information to future nodes, ultimately transforming inputs to outputs by the end of the network.
- **Deep learning**, as pictured below, leads to a layering of groups of nodes, or essentially various levels of interconnected neural networks.

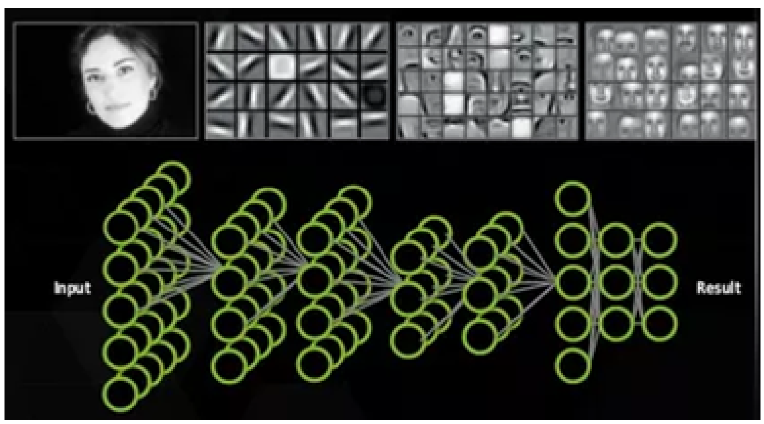

Machine learning leads to the generation of **neural networks** that provide a procedure to analyze presented data.

- **Neural networks** function similar to the system of neurons within a brain—numerous **nodes**, the machine-learning equivalent of neurons, accept, process and rely information to future nodes, ultimately transforming inputs to outputs by the end of the network.
- **Deep learning** leads to a layering of groups of nodes, or essentially various levels of interconnected neural networks.

The process of machine learning can be separated into the dichotomy of **training** and **inference.**

- **Training**: the educational phase of AI, in which a neural network is developed through analyzing massive amounts of historical data.
- **Inference**: the operational phase of AI, in which the neural network developed from training is used for real-life applications to try and estimate desired outputs while potentially having to deal with limited inputs (in contrast to data from the training phase).

Artificial Intelligence

Hardware Needs

•Computers have traditionally relied heavily on **central processing units (CPUs)**, which are the largest, most complex processors that excel in handling singular, intensive tasks but lack in speed of serial computing.

• Recent trends have shifted focus to integrated **accelerators**, or specialized processors that allow parallel computing or optimized data flow to increase speed and efficiency. Some common accelerators include:

- **Co-processors:** processors that provide an extension of traditional CPUs by handling offloaded tasks to lessen the demand on the CPUs.
- **Graphics Processing Units (GPUs):** units that easily handle multiple simultaneous operations, similar to the process of managing numerous pixels for rending images for which they were originally created.
- **Field Programmable Gate Arrays (FPGAs):** reconfigurable programmable logic devices that configure the top layer of systems, helping to delegate procedures across units.
- **Tensor Processing Units (TPUs):** proprietary units developed by Google that allow for great computation speed due to a lesser need for computational precision.
- **Data Processing Units (DPUs):** units that optimize the flow of data between other processors – like CPUs and GPUs – to prevent computational bottlenecks.

Common artificial intelligence applications include:

- **Autonomous driving**: Through a mix of perception, localization and planning utilizing an array of sensors, trained deep neural networks can soon lead to fully autonomous cars with the ability of communication between vehicles.
- **Healthcare**: Tasks such as reading X-rays, MRIs and ultrasound scans can performed by AI, after training with the vast amount of health records and published research.
- **Virtual assistants and chatbots**: Assistants, such as Siri for Apple devices or Alexa for Amazon devices, can provide intelligent suggestions based on recurrent neural networks (RNNs) taking into account prior demands.
- **Image recognition and machine vision**: Systems can be trained in either image classification or object detection, both of which have applications in various fields.
- **Finance**: AI can be implemented for most anything from trading algorithms to detecting fraud to providing financial advice through “robo advisors.”
- **Search engine optimization**: A user’s historical searches can help train neural networks to help improve future accuracy, a process that can be perpetually iterated.

Artificial Intelligence

Quantifying the Opportunity – Total Addressable Market

- The total addressable market (TAM) for AI-related semiconductors is projected to grow immensely at a 63% CAGR through 2020 as firms ramp-up adoption as more applications and potential ROI of AI are realized. A major tailwind for this growth is the further integration of accelerators – predicted to grow at a 94% CAGR – due to inclusion in system architectures to increase efficiency.

- By 2020, we forecast that ~20% of CPUs will be dedicated to AI-related tasks.

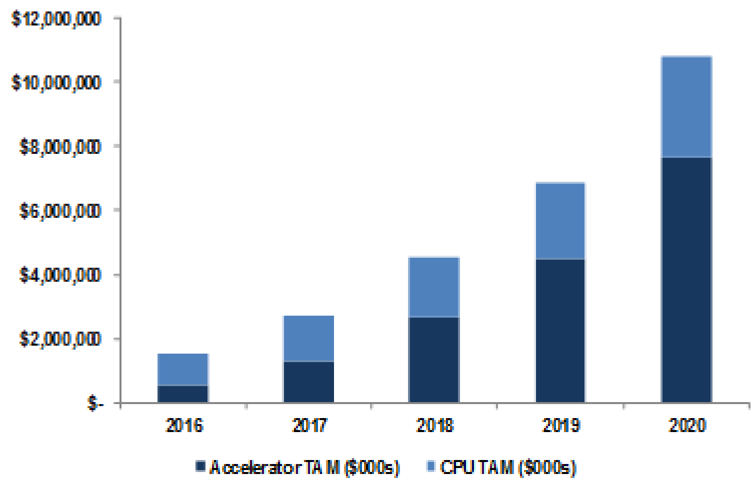

	2015	2016	2017	2018	2019	2020	16-'20 CAGR
CPU TAM ($000s)	$ 274,924	$ 976,122	$ 1,398,300	$ 1,861,275	$ 2,388,672	$ 3,135,000	33.9%
Accelerator TAM ($000s)	$ 66,167	$ 543,152	$ 1,251,188	$ 2,654,457	$ 4,471,586	$ 7,676,630	93.9%
Total ($000s)	$ 341,091	$ 1,519,274	$ 2,649,488	$ 4,515,732	$ 6,860,258	$ 10,811,630	63.3%

The total addressable market (TAM) for AI-related semiconductors is projected to grow immensely at a 63% CAGR through 2020 as firms ramp-up adoption as more applications and potential ROI of AI are realized. A major tailwind for this growth is the further integration of accelerators—predicted to grow at a 94% CAGR—due to inclusion in system architectures to increase efficiency. By 2020, we forecast that ~20% of CPUs will be dedicated to AI-related tasks.

Currently, the most common, currently-accessible accelerators are GPUs due to their ecosystem of development tools and efficient parallel processing capabilities. However, as later systems are developed with AI specifically in mind, accelerators such as co-processors, FPGAs and application-specific integrated circuits (AISCs) are likely garner more share. We believe ASICs remain 1-3 years away from being economically viable, while FPGAs are likely relegated to inferencing applications which remain early in their ramp.

We believe processing players are currently acquiring and developing assets to address the market with major players including NVDA, INTC, AMD, XLNX and QCOM all looking to position themselves in the burgeoning market. We believe NVDA was the first to recognize and adjust its portfolio to address the market's potential, and therefore has become a thought leader in the space. However, other players have taken notice and we expect FPGAs from INTC and XLNX, GPUs from AMD, co-processors from INTC, and ASICs from various players to become increasingly competitive.

Emerging Technologies

100G

Increasingly fast connections are being demanded and the market is beginning to shift directly from 10G to 100G, and even 400G.

100G is beginning to see adoption in core and metro transport, as well as for box-to-box interconnect over short distances.

100G Semiconductor Competitors

Broadcom	Semtech
Marvell	Microsemi
MA-COM	Inphi
MaxLinear	Intel

Networks are increasingly transitioning to 100G to keep up with exploding data and bandwidth requirements. The core of the network is the first to be upgraded, and in many cases, networks are bypassing 40G and directly upgrading to 100G from 10G. We believe this 100G upgrade cycle is a multi-year trend in the industry.

Today, 100G is beginning to see use in core and metro transport, as well as for router-to-router or switch-to-switch interconnect over short distances, often within datacenters.

100G Semiconductor Competitors: Prominent suppliers to the space include Broadcom, Marvell, MA-COM, MaxLinear, Semtech, Microsemi, Inphi, and Intel. ASICs and FPGAs will likely also play an important role as the market gets going.

Emerging Technologies

Antenna Tuning/Control

Global LTE/LTE-A enabled smartphones are being designed with various antenna configurations that support an ever-growing number of cellular frequency bands. An antenna tuner can extend the frequency range and performance of a single antenna to optimize the wavelength of the antenna, match the signal path and dramatically improve wireless radio performance. The tuner makes for one antenna that behaves as if it were multiple antennas with perfect signal quality. All else equal, this means a higher data rate, a more stable signal, less searching for a signal and, ultimately, better handset battery life.

Semiconductor Vendors With the Most Direct Benefit

Qorvo	Murata
Skyworks	WiSpry
Qualcomm	Ethertronics

Antenna tuning (antenna control) is a newer, but ramping market that presents a large opportunity for semiconductor vendors as the LTE handset market matures and evolves into LTE-A and 5G. Adding more cellular bands to a handset requires components that support wider bandwidths with enhanced signal linearity. With various antenna configurations, the antenna signal path is becoming increasingly difficult to navigate. An antenna tuner can, in theory, extend the frequency range and performance of a single antenna to optimize the wavelength of the antenna, match the signal path and dramatically improve wireless radio performance. The tuner makes for one antenna that behaves as if it were multiple antennas with perfect signal quality. All else equal, this means a higher data rate, a more stable signal, less searching for a signal and, ultimately, better battery life.

As we add more frequency bands of LTE spanning a significantly wider bandwidth, the need has emerged for multiple antennas to address these bands, which today means a primary antenna and a diversity antenna. Each OEM architecture is ultimately different, but because device size continues to decrease and the physical space allowable to the RF Front End is shrinking/stable, today's handsets cannot support an ever-higher number of antennas. These are clearly divergent and unsustainable trends. The antenna tuner can dynamically match one antenna with multiple different band frequencies, thereby limiting the need for more antennas (save on cost and size) while increasing the antenna efficiency by more quickly aligning with the network signal (save on battery life). This is a dual benefit and is functionally cheaper than adding an additional antenna to the device.

Emerging Technologies

Picocells and Femtocells (Small Cells)

Picocells are smaller versions of base stations, used in cellular networks to densify network coverage, or provide additional coverage over a small area. Picocells are usually used to extend coverage to indoor areas where outdoor signals do not reach well, or to add network capacity in areas with very dense phone usage, such as train stations. 5G is expected to drive small cell adoption as well to fill coverage gaps.

Femtocells provide smaller, more focused coverage than picocells, targeting consumers. Femtocells are installed by home users to get better indoor voice and data coverage, supporting up to 5 voice/data connections.

Pico/Femto Semiconductor Competitors

Broadcom
Qualcomm
NXP
Marvell/Cavium
Intel
ADI

Picocells are small versions of cellular base stations, ranging in size from a laptop computer to a suitcase. Picocells historically have been used primarily in cellular networks as a means of extending coverage. Besides providing a wireless signal in tough-to-reach pockets, picocells also are frequently used to add voice and data capacity—a capability that repeaters and antenna systems cannot satisfy. Picocells will support multiple wired and wireless backhaul technologies and can help smooth the transition for wireless carriers as they migrate to more advanced networks (LTE, LTE-Advanced, and 5G). The most suitable market is comprised of establishments with between 20 and 100 employees.

Femtocells are the precursor to picocells in that they are intended to serve a smaller group of users (2-5 users). They are most commonly utilized in next-generation network infrastructure technologies such as LTE-A. Deployments among femtocell vendors have been lower than anticipated over the early years as the focus has remained primarily on the buildout of macro basestations.

Semiconductor Competitors: Qualcomm, Broadcom, Intel/, Marvell/Cavium, NXP, and ADI represent the companies with the greatest presence within this relatively nascent space.

Emerging Technologies

Wearables

The highly-publicized Apple Watch is possibly the most recognizable wearable to hit the market so far. The market for "wearables" and the broader "Internet of Things" market are redefining how the world communicates and how we interact with the technology in our lives. Often wearables serve as ancillary devices tethered to a mobile handset or PC. We won't go so far as to size the market given the variety of definitions that "wearables" and "IoT" encompasses, but we do believe both will present a meaningful opportunity for mobile. Common components in wearables include application processors, wireless connectivity, sensing, and low-density memory.

Semiconductor Vendors Poised to Benefit from Growth of Wearables

Broadcom	Maxim
Qualcomm	Texas Instruments
Intel	Silicon Labs
Skyworks	Microchip
Qorvo	Cypress

The proliferation of mobile computing is beginning to extend beyond the reach of just smartphones and tablets. An increasing number of mobile device manufacturers have introduced a variety of wearable computing technology over the past several years. Our fundamental view on the wearables market is that it will be dominated by a handful of products from leading OEMs, while the rest of the market will remain highly fragmented and mainly cater to niche verticals. In the long run, we believe wearables will become an essential part of the always-on, always-connected lifestyle. They likely will have a profound impact across industries. The collection, analysis, and recommendations on currently uncollected, unsorted user data will ultimately drive the potential of wearables.

Wearable Semiconductor Competitors: The same wireless semiconductor vendors that lead in smartphone and tablet platforms are likely to benefit from the application-processing and wireless-connectivity demands of wearable computing platforms. Companies like Qualcomm, Broadcom, Intel, Skyworks, and Qorvo, to name a few, are likely to be at the forefront of the trend. Many sensing and processing companies, like Maxim, Texas Instruments, Silicon Labs, Microchip, and Cypress, are also likely to benefit.

Emerging Technologies

LoRa – Industrial IoT Connectivity

The Industrial IoT is a fast growing market as cloud service providers, enterprises, and municipalities seek to become "smarter." Gartner projects that by 2020 there will be 20.4 billion IoT end points, almost all of which will require connectivity.

In most of these instances, power consumption and range are critical; however, data rates/bandwidth are less important. While many of these connections can be served by common wireless technologies such as WiFi, Bluetooth, and LTE, LoRa is an emerging and compelling play on the industrial IoT, offering a combination of low power consumption and long-range connectivity. Key applications include use cases where sensors are operated by batteries that need to last for up to 10+ years. LoRa offers up to 50Kbps uplink at ranges up to 15KM.

LoRa is currently being evangelized by a 500+ member alliance led by Semtech and including Cisco, Alibaba, IBM, Gemalto, Renesas, Orange, and STMicro amongst others. Cloud service providers have begun deploying LoRa for enterprise and commercial uses including smart metering, asset tracking, and environmental monitoring.

The Industrial Internet of Things is a fast-growing market as cloud service providers, enterprises, and municipalities seek to become "smarter." Gartner projects that by 2020, there will be 20.4 billion IoT end points, nearly all of which will require connectivity. In most of these instances, power consumption and range are critical; however, data rates/bandwidth are less important. While many of these connections can be served by common wireless technologies such as WiFi, Bluetooth, and LTE, LoRa (Long Range) is an emerging and compelling play on the industrial IoT, offering a combination of low power consumption and long-range connectivity. LoRa technology is owned by Semtech which formed the LoRa Alliance, with the goal of developing protocols and an open media access control layer for LoRa in order to ease interoperability between devices. LoRa operates at 850MHz -1GHz using unlicensed spectrum making it available to anyone.

Key applications include use cases where sensors are operated by batteries that need to last for up to 10+ years. LoRa offers up to 50Kbps uplink at ranges up to 15KM. LoRa is currently being evangelized by a 500+ member alliance led by Semtech and including Cisco, Alibaba, IBM, Gemalto, Renesas, Orange, and STMicro, among others. Cloud service providers have begun deploying LoRa for enterprise and commercial uses including smart metering, asset tracking, and environmental monitoring.

For More Information…

Questions?
Comments?
Suggestions?

Rick Schafer
(720) 554-1119
rick.schafer@opco.com

Joshua Buchalter
(212) 667-8387
joshua.buchalter@opco.com

We hope you enjoyed reading our *Semiconductor Primer 10.0* and welcome your questions, comments, and suggestions.

Some of the forecast data in this report is available in Excel format.

Stock Prices of Other Companies in this Report as of 12/4/17

Acer, Inc. (2353-TAI, NT$20.15, Not Covered)
ASML Holding NV (ASML-AMS, €145.55, Not Covered)
ASUSTek Computer, Inc. (2357-TAI, NT$268.5, Not Covered)
Canon, Inc. (7751-JP, ¥4290, Not Covered)
China Mobile Ltd. (941-HKG, HK$78.35, Not Covered)
China Telecom Corp. Ltd. (728-HKG, HK$3.8600001, Not Covered)
China Unicom (Hong Kong) Ltd. (762-HKG, HK$11.2, Not Covered)
Compal Electronics, Inc. (2324-TAI, NT$20.95, Not Covered)
Cree, Inc. (CREE, $35.08, Not Covered)
Cypress Semiconductor Corp. (CY-NASDAQ, $14.965, Not Covered)
Gigabyte Technology Co., Ltd. (2376-TAI, NT$46.35, Not Covered)
HP, Inc. (HPQ, $21.1338, Not Covered)
HTC Corp. (2498-TAI, NT$68.1, Not Covered)
Infineon Technologies AG (IFX-ETR, €22.725, Not Covered)
Inphi Corp. (IPHI, $40.0568, Not Covered)
Integrated Device Technology, Inc. (IDTI-NASDAQ, $29.86, Not Covered)
International Business Machines Corp. (IBM, $155.42, Not Covered)
Lenovo Group Ltd. (992-HKG, HK$4.4, Not Covered)
LG Electronics, Inc. (066570-KRX, W94300, Not Covered)
MACOM Technology Solutions Holdings, Inc. (MTSI, $32.18, Not Covered)
MaxLinear, Inc. (MXL - NYS, $25.54, Not Covered)
MediaTek, Inc. (2454-TAI, NT$318, Not Covered)
Mellanox Technologies Ltd. (MLNX-NASDAQ, $58.55, Not Covered)
Microchip Technology, Inc. (MCHP-NASDAQ, $87.83, Not Covered)
Micron Technology, Inc. (MU-NASDAQ, $41.2975, Not Covered)
Murata Manufacturing Co. Ltd. (6981-TKS, ¥15020, Not Covered)
Nanya Technology Corp. (2408-TAI, NT$73.8, Not Covered)
ON Semiconductor Corp. (ON - NASDAQ, $19.51, Not Covered)
Pegatron Corp. (4938-TAI - TAI, NT$69.7, Not Covered)
Quanta Computer, Inc. (2382-TAI - TAI, NT$60.8, Not Covered)
Realtek Semiconductor Corp. (2379-TAI - TAI, NT$107, Not Covered)
ROHM Co., Ltd. (6963-JP - TKS, ¥11250, Not Covered)
Samsung Electronics Co., Ltd. (005930-KRX, W2563000, Not Covered)
Seagate Technology Plc (STX - NASDAQ, $39.75, Not Covered)
Silicon Laboratories, Inc. (SLAB - NASDAQ, $88.075, Not Covered)
SK hynix, Inc. (000660-KRX - KRX, W77900, Not Covered)
Sony Corp. (SNE - NYSE, $46.02, Not Covered)
STMicroelectronics NV (STM - NYSE, $22.345, Not Covered)
Taiwan Semiconductor Manufacturing Co., Ltd. (2330-TAI - TAI, NT$229.5, Not Covered)
Toshiba Corp. (6502-TKS - TKS, ¥280, Not Covered)
United Microelectronics Corp. (2303-TAI - TAI, NT$15.55, Not Covered)
Vanguard International Semiconductor Corp. (5347-TAI - ROCO, NT$68.7, Not Covered)
Western Digital Corp. (WDC - NASDAQ, $77.65, Not Covered)
Xilinx, Inc. (XLNX - NASDAQ, $69, Not Covered)
ZTE Corp. (763-HKG - HKG, HK$27.15, Not Covered)

Disclosure Appendix

Oppenheimer & Co. Inc. does and seeks to do business with companies covered in its research reports. As a result, investors should be aware that the firm may have a conflict of interest that could affect the objectivity of this report. Investors should consider this report as only a single factor in making their investment decision.

Analyst Certification - The author certifies that this research report accurately states his/her personal views about the subject securities, which are reflected in the ratings as well as in the substance of this report. The author certifies that no part of his/her compensation was, is, or will be directly or indirectly related to the specific recommendations or views contained in this research report.

Potential Conflicts of Interest:

Equity research analysts employed by Oppenheimer & Co. Inc. are compensated from revenues generated by the firm including the Oppenheimer & Co. Inc. Investment Banking Department. Research analysts do not receive compensation based upon revenues from specific investment banking transactions. Oppenheimer & Co. Inc. generally prohibits any research analyst and any member of his or her household from executing trades in the securities of a company that such research analyst covers. Additionally, Oppenheimer & Co. Inc. generally prohibits any research analyst from serving as an officer, director or advisory board member of a company that such analyst covers. In addition to 1% ownership positions in covered companies that are required to be specifically disclosed in this report, Oppenheimer & Co. Inc. may have a long position of less than 1% or a short position or deal as principal in the securities discussed herein, related securities or in options, futures or other derivative instruments based thereon. Recipients of this report are advised that any or all of the foregoing arrangements, as well as more specific disclosures set forth below, may at times give rise to potential conflicts of interest.

Important Disclosure Footnotes for Companies Mentioned in this Report that Are Covered by Oppenheimer & Co. Inc:

Stock Prices as of December 7, 2017

Apple Inc. (AAPL - NASDAQ, $169.01, PERFORM)
Analog Devices (ADI - NASDAQ, $85.41, OUTPERFORM)
Advanced Micro Devices (AMD - NYSE, $10.00, PERFORM)
Amazon.Com, Inc. (AMZN - NASDAQ, $1,152.35, OUTPERFORM)
Broadcom Ltd. (AVGO - NYSE, $263.89, OUTPERFORM)
Best Buy Co. Inc. (BBY - NYSE, $60.97, PERFORM)
Cavium, Inc. (CAVM - NASDAQ, $84.60, PERFORM)
CEVA Inc. (CEVA - NASDAQ, $44.80, OUTPERFORM)
Comcast (CMCSA - NASDAQ, $38.63, OUTPERFORM)
Cree, Inc. (CREE - NASDAQ, $34.35, PERFORM)
Cirrus Logic (CRUS - NASDAQ, $49.71, PERFORM)
Cisco Systems (CSCO - NASDAQ, $37.41, OUTPERFORM)
Facebook, Inc. (FB - NASDAQ, $176.06, OUTPERFORM)
F5 Networks (FFIV - NASDAQ, $131.47, PERFORM)
Alphabet Inc. (GOOG - NASDAQ, $1,018.38, OUTPERFORM)
International Business Machines (IBM - NYSE, $154.10, PERFORM)
Intel Corp. (INTC - NASDAQ, $43.45, PERFORM)
Juniper Networks, Inc. (JNPR - NASDAQ, $27.96, PERFORM)
Monolithic Power Systems (MPWR - NASDAQ, $111.24, OUTPERFORM)
Marvell Technology Group (MRVL - NASDAQ, $21.82, OUTPERFORM)
Microsemi Corp. (MSCC - NASDAQ, $51.10, PERFORM)
Microsoft Corporation (MSFT - NASDAQ, $82.78, OUTPERFORM)
Maxim Integrated Products (MXIM - NASDAQ, $51.50, PERFORM)
Netflix, Inc. (NFLX - NASDAQ, $185.30, OUTPERFORM)

NetApp, Inc. (NTAP - NASDAQ, $57.21, PERFORM)
NVIDIA Corp. (NVDA - NASDAQ, $189.26, PERFORM)
NXP Semiconductors NV (NXPI - NASDAQ, $114.33, PERFORM)
Universal Display Corp. (OLED - NASDAQ, $164.05, PERFORM)
Palo Alto Networks Inc. (PANW - NYSE, $141.37, OUTPERFORM)
QUALCOMM Incorporated (QCOM - NASDAQ, $64.98, PERFORM)
Qorvo, Inc. (QRVO - NASDAQ, $70.10, PERFORM)
Sprint (S - NYSE, $5.69, PERFORM)
Semtech Corp. (SMTC - OTC, $32.50, OUTPERFORM)
Skyworks Solutions, Inc. (SWKS - NASDAQ, $96.47, OUTPERFORM)
AT&T, Inc. (T - NYSE, $36.11, PERFORM)
T-Mobile (TMUS - NYSE, $62.12, OUTPERFORM)
Texas Instruments (TXN - NYSE, $97.02, OUTPERFORM)
VMware, Inc. (VMW - NYSE, $116.67, OUTPERFORM)
Verizon (VZ - NYSE, $50.68, PERFORM)

All price targets displayed in the chart above are for a 12- to- 18-month period. Prior to March 30, 2004, Oppenheimer & Co. Inc. used 6-, 12-, 12- to 18-, and 12- to 24-month price targets and ranges. For more information about target price histories, please write to Oppenheimer & Co. Inc., 85 Broad Street, New York, NY 10004, Attention: Equity Research Department, Business Manager.

Oppenheimer & Co. Inc. Rating System as of January 14th, 2008:

Outperform(O) - Stock expected to outperform the S&P 500 within the next 12-18 months.

Perform (P) - Stock expected to perform in line with the S&P 500 within the next 12-18 months.

Underperform (U) - Stock expected to underperform the S&P 500 within the next 12-18 months.

Not Rated (NR) - Oppenheimer & Co. Inc. does not maintain coverage of the stock or is restricted from doing so due to a potential conflict of interest.

Oppenheimer & Co. Inc. Rating System prior to January 14th, 2008:

Buy - anticipates appreciation of 10% or more within the next 12 months, and/or a total return of 10% including dividend payments, and/or the ability of the shares to perform better than the leading stock market averages or stocks within its particular industry sector.

Neutral - anticipates that the shares will trade at or near their current price and generally in line with the leading market averages due to a perceived absence of strong dynamics that would cause volatility either to the upside or downside, and/or will perform less well than higher rated companies within its peer group. Our readers should be aware that when a rating change occurs to Neutral from Buy, aggressive trading accounts might decide to liquidate their positions to employ the funds elsewhere.

Sell - anticipates that the shares will depreciate 10% or more in price within the next 12 months, due to fundamental weakness perceived in the company or for valuation reasons, or are expected to perform significantly worse than equities within the peer group.

Distribution of Ratings/IB Services Firmwide

			IB Serv/Past 12 Mos.	
Rating	Count	Percent	Count	Percent
OUTPERFORM [O]	314	59.02	131	41.72
PERFORM [P]	213	40.04	68	31.92
UNDERPERFORM [U]	5	0.94	4	80.00

Although the investment recommendations within the three-tiered, relative stock rating system utilized by Oppenheimer & Co. Inc. do not correlate to buy, hold and sell recommendations, for the purposes of complying with FINRA rules, Oppenheimer & Co. Inc. has assigned buy ratings to securities rated Outperform, hold ratings to securities rated Perform, and sell ratings to securities rated Underperform.

Note: Stocks trading under $5 can be considered speculative and appropriate for risk tolerant investors.

Company Specific Disclosures

Oppenheimer & Co. Inc. makes a market in the securities of AAPL, ALTR, AMZN, ASML, AVGO, BRCD, CEVA, CMCSA, CREE, CRUS, CSCO, CY, DCM, ERIC, EXAR, FB, FFIV, GOOG, IDTI, INTC, ISSI, JNPR, MPWR, MRVL, MSCC, MSFT, MXIM, NFLX, NTAP, NVDA, NXPI, OLED, PMCS, QCOM, QRVO, SIGM, SMTC, STX, SWKS, T, TMUS, TXCC, WDC, XLNX, MCHP and ON.

Oppenheimer & Co. Inc. expects to receive or intends to seek compensation for investment banking services in the next 3 months from ANAD, CMCSA, CREE, FB, IDTI, IPHI, ISSI, MSCC and OLED.

Oppenheimer & Co. Inc. and its affiliates, in the aggregate, beneficially own 1% or more of a class of equity securities issued by CY.

Additional Information Available

Please log on to http://www.opco.com or write to Oppenheimer & Co. Inc., 85 Broad Street, New York, NY 10004, Attention: Equity Research Department, Business Manager.

Other Disclosures

This report is issued and approved for distribution by Oppenheimer & Co. Inc. Oppenheimer & Co. Inc. transacts business on all principal exchanges and is a member of SIPC. This report is provided, for informational purposes only, to institutional and retail investor clients of Oppenheimer & Co. Inc. and does not constitute an offer or solicitation to buy or sell any securities discussed herein in any jurisdiction where such offer or solicitation would be prohibited. The securities mentioned in this report may not be suitable for all types of investors. This report does not take into account the investment objectives, financial situation or specific needs of any particular client of Oppenheimer & Co. Inc. Recipients should consider this report as only a single factor in making an investment decision and should not rely solely on investment recommendations contained herein, if any, as a substitution for the exercise of independent judgment of the merits and risks of investments. The analyst writing the report is not a person or company with actual, implied or apparent authority to act on behalf of any issuer mentioned in the report. Before making an investment decision with respect to any security recommended in this report, the recipient should consider whether such recommendation is appropriate given the recipient's particular investment needs, objectives and financial circumstances. We recommend that investors independently evaluate particular investments and strategies, and encourage investors to seek the advice of a financial advisor. Oppenheimer & Co. Inc. will not treat non-client recipients as its clients solely by virtue of their receiving this report. Past performance is not a guarantee of future results, and no representation or warranty, express or implied, is made regarding future performance of any security mentioned in this report. The price of the securities mentioned in this report and the income they produce may fluctuate and/or be adversely affected by exchange rates, and investors may realize losses on investments in such securities, including the loss of investment principal. Oppenheimer & Co. Inc. accepts no liability for any loss arising from the use of information contained in this report, except to the extent that liability may arise under specific statutes or regulations applicable to Oppenheimer & Co. Inc. All information, opinions and statistical data contained in this report were obtained or derived from public sources believed to be reliable, but Oppenheimer & Co. Inc. does not represent that any such information, opinion or statistical data is accurate or complete (with the exception of information contained in the Important Disclosures section of this report provided by Oppenheimer & Co. Inc. or individual research analysts), and they should not be relied upon as such. All estimates, opinions and recommendations expressed herein constitute judgments as of the date of this report and are subject to change without notice. Nothing in this report constitutes legal, accounting or tax advice. Since the levels and bases of taxation can change, any reference in this report to the impact of taxation should not be construed as offering tax advice on the tax consequences of investments. As with any investment having potential tax implications, clients should consult with their own independent tax adviser. This report may provide addresses of, or contain hyperlinks to, Internet web sites. Oppenheimer & Co. Inc. has not reviewed the linked Internet web site of any third party and takes no responsibility for the contents thereof. Each such address or hyperlink is provided solely for the recipient's convenience and information, and the content of linked third party web sites is not in any way incorporated into this document. Recipients who choose to access such third-party web sites or follow such hyperlinks do so at their own risk.

This research is distributed in the UK and elsewhere throughout Europe, as third party research by Oppenheimer Europe Ltd, which is authorized and regulated by the Financial Conduct Authority (FCA). This research is for information purposes only and is not to be construed