

IBM

READ ONLY STORAGE

CN						ADR		W REGISTER							X REG						
P	0	1	2	3	4	5	P	LP	P	1	8	4	2	1	P	8	4	2			
									P	3	4	5	6	7	P	0	1	2			
SA				CH			CL			CA				CB		CM		CU			
P	0	1	2	3	0	1	2	3	A	0	1	2	3	0	1	0	1	2	0	1	A
CR		CD				CF				CG		CV		CC							
P		0	1	2	3	0	1	2	0	1	0	1	0	1	0	1	2				

COUNT REGISTER

P	8	4	2	1	8	4	2	1	P	8	4	2	1	8	4	2	1
P	0	1	2	3	4	5	6	7	P	0	1	2	3	4	5	6	7

CHANNEL NUMBER ONE

DATA REGISTER								KEY				COMMAND					
P	8	4	2	1	8	4	2	1	P	8	4	2	1	8	4	2	1
P	0	1	2	3	4	5	6	7	P	0	1	2	3	4	5	6	7

THE IBM SYSTEM/360

BY RICHARD S. TEDLOW

A LOOK BACK AT THE CREATION OF A COMPUTING HISTORY GIANT



the United States Air Force in order to detect hostile aircraft. The RAYDAC (Raytheon Digital Automatic Computer), to choose another example, was built by Raytheon for the Naval Air Missile Test Center in California.

To be sure, not all computers were one-offs. IBM marketed both a scientific line and a business line of processors in the late 1950s. However, as their names suggest, one line was targeted at business and the other at scientific markets. Businesses tended to want computers for simple calculations involving text and decimal numbers, performed at a reasonable speed. Computers targeting the scientific market, in contrast, typically had to be able to perform very sophisticated calculations on immense data sets.

Perhaps the single most important characteristic of *all* computers prior to the 360 was that they could not communicate with each other. Born in the midst of World War II, during which the need to deal with huge amounts of data was paramount, analysis rather than communication was what computers were developed to deliver.

With the remarkable progress in computers from vacuum tubes to transistors after World War II, it became progressively more obvious that computer users wanted computers that were compatible—that could use the same software and peripherals as their businesses grew. Compatibility became the “Holy Grail” of the industry in the late 1950s and 1960s.

Compatibility was important for any number of reasons. Without it, every time customers made a change from one central processing unit (the hardware, which was the heart of the computing system) to another, they had to change everything else. All the peripherals—printers, input/output devices, magnetic storage devices, and so on—had to be changed as well. This was no small matter. Some 44 peripheral devices were announced that fateful day in the early spring of 1964 along with the System/360 CPU itself.⁴

Hardware was only half the problem that had to be solved for compatibility to become a reality. The other half, perhaps even more difficult to manage, was software. One of the earliest theoretical works about software was published by the brilliant Alan Turing in an essay written in 1935.⁵ Software was just coming into its own as a discrete field in the 1950s. *Wikipedia* defines software as a set of instruc-

The IBM System/360 is rightly viewed as one of the great new product introductions in the history of business. The 360 transformed both the company that introduced it and the industry of which that company was a part. Not incidentally, the struggle to create the new world dominated by the 360 transformed the lives of many of the key people involved in what was to become a great drama.

Thomas J. Watson Jr., CEO and chairman of the board, announced the System/360 on April 7, 1964, a half century ago. There was no hiding of lights under bushels on the occasion of the announcement. The company hired a special train to take about 200 reporters from Grand Central Station to its facilities in Poughkeepsie, New York, where Watson spoke; and press conferences were held in 165 cities across the United States and in 14 other countries.¹

“We are not at all humble,” began Watson, “about asking you to come here today . . . to share with us the most important product announcement that this corporation has ever made in its history.”² This was not an overstatement. Many years later, Andy Grove remarked that the 360 casted a shadow that lasted decades.³

Realizing a Computer of the Future

What was it about the 360 that made it an artifact of such supreme importance? Central to the 360 were two related ideas. One was that the “computer of the future” should be able to serve the needs of business, government, and science. During the 1950s, computers were often special-purpose machines. IBM, for example, built the SAGE (Semi-Automatic Ground Environment) computer system for

Preceding page:
Close-up detail of IBM
System/360 Model 30.

¹ Chuck Boyer, *The 360 Revolution* (Armonk, NY: IBM Corporation, 2004).

² Quotation courtesy of the IBM Archives.

³ In conversation with the author, 2005.

⁴ Emerson W. Pugh, Lyle R. Johnson, and John H. Palmer, *IBM's 360 and Early 370 Systems* (Cambridge: MIT Press, 1991), 167.

⁵ “History of Software,” *Wikipedia, The Free Encyclopedia*

⁶ Wikipedia contributors, “Software,” *Wikipedia, The Free Encyclopedia*, <http://en.wikipedia.org/wiki/Software>.

⁷ See <http://www-03.ibm.com/ibm/history/ibm100/us/en/>

tions that tells computer hardware what to do, a deceptively simple description for what can be a phenomenon of great complexity.⁶

In the 1950s, software for one computer was not portable to another. Thus, every time a customer wanted to change computers, they would have to write a whole new set of program instructions for the new machine. This was time-consuming, intricate work in which mistakes were not uncommon.

As a result, each decision to change a computer required the expenditure of time, effort, energy, and money. If you wanted a more powerful machine because you were growing or you wanted an additional machine because you were expanding geographically, new software and peripherals were required.

For the vendor, this situation was highly problematic. It led to product proliferation, which made it impossible to capture scale economies. It meant you could not achieve customer “lock-in.” That is to say, each time a customer wanted to change computers, he owed it to himself to survey the whole competi-

tive landscape rather than to stick with you because of the necessity of rewriting software and buying new peripherals. Incompatibility was the enemy of brand loyalty.

A Threat to Service

For IBM, incompatibility was a problem of special urgency. Founded in 1911, IBM was a company built on service. Through the electromechanical era into the electronic era, IBM’s service was always beyond compare. Yet by 1960, IBM’s capability to provide superior service was being severely tested. The technological accomplishments of the 360 are so striking that marketing, sales, and service can be easily overlooked. Yet they were critically important, especially for IBM.

IBM has always been focused first and foremost on sales and service. The man who imprinted his approach to business indelibly on the firm was Thomas J. Watson Sr., CEO from 1914 to 1956. He was pre-eminently a salesman who used to tell his engineers that IBM did not sell punched card ma-

IBM President and CEO Thomas J. Watson Jr. (left) and IBM Senior Vice President Vin Learson with IBM System/360 Model 20, ca. 1966.





Typical IBM System/360
Model 65 installation, 1965.

chines but rather “a service that satisfies.”⁷ The eight CEOs who succeeded Watson all came up through sales and marketing. None was a technologist.

In the words of his son and heir, Tom Jr., “The main aim of our business is service, to help the customer solve his problems no matter how many problems this may create for us.” Young Tom also said that IBM hopefully delivered “cutting edge equipment, hopefully all sorts of pioneering efforts, hopefully Nobel Prizes . . . but the service is something most companies forget.”⁸

Incompatibility was threatening IBM’s distinctive competitive advantage in service. By the early 1960s, according to Watson, “our product line had become wildly disorganized.” For each processor, such as the 1401, 1620, 7070, 7080, and 7090, IBM was commit-

ted to providing its customers with programming, training, and field service.⁹ IBM had to train its own people to perform these tasks before they could educate others. The company could not move a field service person from a 1401 to a 7090 if demand required because the codes and the peripherals were different. That service person would have to be retrained.

It is in the marketing function that we can find the answer to one of the great puzzles of the 360. That puzzle is: Why did IBM, the industry leader by a wide margin, take this giant leap into the unknown? It is rare enough for any company to make a bet of this magnitude. It is close to unheard of for such a bet to be made by a company that was winning the game as it was being played.

Despite the real problems caused by product proliferation that incompatibility generated, IBM was by far the dominant force in the industry. In 1960, its market share was about 75 percent. The competitors, some of which were big companies loaded with engineering expertise, included General Electric (GE), Radio Corporation of America (RCA), and American Telephone and Telegraph Company (AT&T), who were dwarfs by comparison to IBM in computing. Nearly half of all computers in the world by the early 1960s were 1401-type systems. But in order to maintain its dominance, the company would have to bring out a new line of computers that was not compatible with any of the machines it was marketing in 1960. It risked sacrificing its leadership and letting everybody start all over again on an equal footing.

The Series That Wasn't: Brooks vs. Evans

In 1960 and 1961, IBM's 7000 series was due for a refresh. The 8000 series was being created to fill this role. The 8000 represented an incremental improvement over its predecessor. Fred Brooks, a 29-year-old Harvard PhD and a computer genius, was the product champion.

The specs and the pricing had all been worked out. There was even a prototype. Everyone was happy with the presentation Brooks made about the device in January of 1961 . . . with one exception. The exception was to matter a great deal.

T. Vincent Learson was a senior vice president of IBM in 1961. He had majored in mathematics at Harvard, from which he graduated in 1935, and went to work at IBM in sales and marketing. He spent his entire career with the company.

The 8000 proposal represented a clear advance over the 7000. In the short- and medium-term future, it would enable IBM to maintain its leadership in computers. But over the long-term, it would intensify the product proliferation problem with which IBM already had to contend.

Don Spaulding, Learson's chief of staff and another major player in what would become the 360 decision, prepared a lengthy memorandum on product policy. He believed that IBM already had

too many different computers requiring too much support and too many peripherals.¹⁰ Spaulding felt the task at hand was to focus and simplify.

Spaulding's memorandum reinforced Learson's own concerns. Another important influence on Learson's thinking was a course on industrial dynamics, which he attended at MIT along with a group of computer users. This experience contributed to his conviction that computer applications would rapidly expand.¹¹ A bold move away from mere record keeping and toward more sophisticated uses was called for. The 8000 may have been an improvement on its predecessor, but it was, nevertheless, a move in the wrong direction.

In 1961, Learson transferred Bob Evans, the manager of processing systems at the General Products Division, from Endicott, New York, to Poughkeepsie to serve as head of planning and development for the Data Systems Division and to help think strategically about charting a new course for overall product policy. Thirty-four years old, Evans had ascended quickly at IBM after joining the firm in 1951. Sleeping four hours a day and working the remaining 20, Evans piloted the 1400 series at the General Products Division in Endicott from the projected 5,000 units to four times that many.

On one December day in 1960, Evans found himself in Milwaukee calling on happy 1401 customers to learn more about what they wanted him

Learson said to him, "Bob, they have this 8000 series up in Poughkeepsie. Go up and look at it. If it's right, do it. If it's not right, do what's right."

to be doing with new products in that family. He received a phone call at about 1:30 in the afternoon informing him that he should be in Learson's office in New York City that evening at eight o'clock. So,

⁹ See www.03.ibm.com/ibm/history/multimedia/wav/ibmservicewav.wav

⁹ Carliss Y. Baldwin and Kim B. Clark, *Design Rules: The Power of Modularity*, Volume 1, (Cambridge: MIT Press, 2000) p. 171.

¹⁰ T.A. Wise, "IBM's \$5,000,000,000 Gamble," *Fortune*, September 1966, 118-21.

¹¹ *Ibid.*

¹² See The IBM System/360 40th Anniversary, Computer History Museum, http://www.youtube.com/watch?v=8c0_Lzb1CJw.

What was the new product that IBM wanted to bring to market going to be? No one knew. But many in the company knew what had to be avoided. Product proliferation had to be conquered.

he canceled his plans and made the trip. Learson said to him, “Bob, they have this 8000 series up in Poughkeepsie. Go up and look at it. If it’s right, do it. If it’s not right, do what’s right.”¹²

This is why Evans was relocated from Endicott to Poughkeepsie. Late in March 1961, he submitted a report to the president of the division concluding that the 8000 series should die. The technology was wrong, and the incompatibility problem was not being addressed.

Forty years later, Fred Brooks remembered vividly the intense emotion of the product fight that went on in 1960 and 1961. He and his team had been optimistic following the presentation in Poughkeepsie in January of 1961. The 8000 was designed to meet an immediate sales need “because the competition was eating us, and we were becoming obsolete.”¹³ The fight over its future continued from January until May.

Twice, the battle was escalated to IBM’s top management. In March 1961, it looked like Brooks had won and the 8000 would proceed. In May, however, the corporate management committee again considered the product roadmap, and this time Evans won.

After he had prevailed in May, Evans well understood that there was no plan in place for the Data Systems Division in Poughkeepsie to do anything.

He invited 25 of the top people at the division to attend an off-site meeting at the Gideon Putnam Hotel in the old resort town of Saratoga Springs, New York. Fred Brooks was among the people present. As he put it, he “just tagged along” to make sure that all his people landed on their feet. He expected to return to research after the meeting.¹⁴

On May 15 at Saratoga Springs, Evans announced to the 8000 team that the product was cancelled. It was his responsibility to reassign those who had championed it to other tasks. Gloom and doom reigned. As one participant said, “There was blood all over the floor.”¹⁵

A New Leader

The first step that Evans took was to develop “temporizers.” If IBM did not make at least incremental improvements to its product line, it would lose share to the competition which, as he put it, “would kick the heck out of IBM’s current products.” Thus, the 7095 was planned to breathe some life into the 7090, and some improvements to the 7080 gave it some life as well. Other temporary improvements were planned and one model was developed in order to satisfy the desire of the sales force for mid-range scientific processing power.¹⁶

“We went back to work hammerin’ and sawin’, and we needed a leader,” said Evans. Evans made what he called the second best decision of his life (the best having been to ask his wife to marry him). That second best decision was to ask Fred Brooks to be the head of the new product launch. He said Brooks “was the best guy around by any measure.”¹⁷

No one was more surprised at this request than Brooks himself. “You could have knocked me over with a feather when he asked me to take the crown jewels of the new plan. That is a big man, and I was absolutely stunned. I talked to one or two senior executives and asked if this was real. Was this something I should do? One of them replied, ‘No one who has ever worked for Bob Evans has regretted it.’”¹⁸ So Brooks accepted the assignment.

What was the new product that IBM wanted to bring to market going to be? No one knew. But many in the company knew what had to be avoided. Product proliferation had to be conquered. IBM was

¹³ Ibid.

¹⁴ Brooks, Fred (Frederick P., Jr.) oral history, September 16, 2007, Oral Histories Online, Lot X4146.2008, Catalog Number 102658255, Computer History Museum.

¹⁵ Wise, “Gamble.”

¹⁶ Bob O. Evans, “System/360: A Retrospective View,” *Annals of the History of Computing*, Vol. 8, No. 2 (April 1986): 163. See http://www.youtube.com/watch?v=8c0_Lzb1CJw.

¹⁷ Ibid.

¹⁸ Ibid.

¹⁹ William O. Ingle and Joseph L. Bower, “The IBM 360: Giant as Entrepreneur,” HBS Case #9-389-003, Rev. April 1, 1998, 7–8.

saying yes to too many good ideas. The result was an overwhelming strain on the company's ability to service customers. The only solution was compatibility.

Indeed, top management wanted a compatible product line well before the technical people—the engineers and designers—knew whether or not this goal could be achieved. In 1959 and 1960, customers were complaining loudly to IBM's sales force about the costs of complexity and incompatibility in IBM's product line.

Unfortunately, IBM's top managers could not command its engineers to do what many thought was technically impossible. How much compatibility could be achieved? Should the instruction set legacies of the 1400 and 7000 machines be perpetuated? Or was it time to start with a clean sheet of paper? These are the questions for which Brooks and his colleagues had to provide answers.

The SPREAD Task Group

During the late summer and early fall of 1961, Watson and Learson initiated discussions with their division heads aimed at defining a new strategy for a new era. The results of these discussions were less than

satisfactory. Learson, therefore, formed a special committee with representatives from all major parts of the company to provide policy guidance.

Learson wanted not only to make the right decision. He wanted to make the decision right. As his vehicle for making the right decision the right way, he impaneled a special task force. Thirteen people were members of what was known as the SPREAD (Systems Programming, Research, Engineering, and Development) task group.

Not a patient man on a good day, Learson, by November, found the progress of this group "hellishly slow." In December, he sequestered these 13 men in the Sheraton New Englander motel in Cos Cob, Connecticut, near Stamford, and told them not to come out until they had reached some conclusions.¹⁹

What emerged on December 28 was their final report. They recommended the introduction of a new product line, which would satisfy the heretofore separate commercial and scientific markets. The new product should be compatible all the way from the basic 1400 machine to the top-of-the-line most sophisticated machine in the 7000 series.

Left: Robert Evans, Vice President, Development, Data Systems Division
Right: Vin Learson, the man who led the System/360 project at the executive level.



Remarkably—incredibly—the task force concluded that since the new product line “must have capabilities not now present in any IBM processor product, the new family of products will not be compatible with our existing processors.”²⁰ Therefore, not only did the task force advocate leapfrogging the industry, their recommendations actually advocated rendering IBM’s present market-dominating products obsolete. It was as if Henry Ford decided in 1921, when the Model T had 55 percent of the automobile market in the United States, to abandon it not only for an enhanced Model T, but for an automobile powered by electricity instead of the internal combustion engine.

The conclusions of the SPREAD report were presented to IBM’s 50 top executives on January 4, 1962. The mellifluous Fred Brooks managed the presentation. The reception was mixed.

The engineers were in favor of it, but other people were shocked. The proposals were met by almost violent opposition from marketing. Their smorgasbord of machines would be done away with. Variety was their ally. That is how they sold. “We have the machine just for you.” With the newly proposed product, customers would have to reprogram, not something salesmen were anxious to tell them. In finance, people were very concerned. Learson saw them slam their books shut. They thought the proposal was too grandiose. The report said, for example, that IBM would spend \$125 million on programming the system at a time when the company spent only about \$10 million a year for programming. As Learson recalled, “The job just looked too big. . . . Everyone recognized it was a gigantic task that would mean all our resources were tied up in one project—and we knew that for a long time we would not be getting anything out of it.”²¹

Nevertheless, Learson concluded, “All right, we’ll do it.” In May 1962, IBM’s corporate management committee formally approved the new product launch.

Gambling on the Unknown

A lot of the drama and a great deal of the uncertainty that surrounded the 360 took place after the decision to move forward with it had been made. IBM did not have a firm understanding of what it was getting itself into. The whole project was almost sunk by the unknown unknowns.

Nothing of the magnitude of the 360 had ever been attempted before—certainly not in this industry. Therefore, IBM had no models to guide it. The

company’s estimates of the cost of hardware and software were so incorrect that the project was in jeopardy on more than one occasion.

The software was a tremendous hurdle. Hundreds of programmers had to write millions of lines of computer code. Nobody had ever tackled that complex of a programming job, and the engineers were under great pressure to get it done.

As late as 1966, a year and a half after the announcement of the 360, Watson, speaking to a group of IBM customers, said, “We are investing nearly as much in System/360 programming as we are in the entire development of System/360 hardware. A few months ago, I was informed that the bill for 1966 was going to be \$40 million. I asked Vin Learson before I left [for this meeting] what he thought the cost would be for 1966, and he said \$50 million. Twenty-four hours later I asked Watts Humphrey, who is in charge of programming production, in the hall here and said, ‘Is this figure about right? Can I use it?’ He said, ‘It’s going to be \$60 million.’ You can see that if I keep asking questions we won’t pay a dividend this year.”²²

Making one big bet on the 360 led to other big bets, which had not been anticipated. Up until the 360, IBM was essentially a company that assembled, marketed, and serviced computers. But for technical reasons, the 360 meant that IBM had to enter component manufacturing, a basic change in the character of the company. In the day of vacuum tubes and transistors, IBM designed the components for circuits, ordered them from other companies (such as Texas Instruments), and then assembled them to its specifications. But with the circuitry required by the 360, those specifications would have to be built into the components from the outset.

As a result, IBM became the world’s largest manufacturer of computer components, an outcome that was neither planned nor welcomed. It also became by far the world’s largest producer of semiconductors. Watson said, “I will never forget how expensive it was to build our first integrated circuit factory. Ordinary plants in those days cost about \$40 per square foot. In the integrated circuit plant, which had to be kept dust free and looked more like a surgical ward than a factory floor, the cost was over \$150. I could hardly believe the bills that were coming through, and I wasn’t the only one who was shocked. The board gave me a terrible time about the costs. ‘Are you really sure you need all this?’ they would say.”²³

IBM's accounting system was inadequate for the manufacturing tasks undertaken. When time came time to close the books for 1965, no one could figure out how much work-in-process inventory the company had. Al Williams, IBM's president, tried to complete his accounting and financial responsibilities in 1965 and couldn't do it. The best estimate he could get for work-in-process inventory was \$150 million, but the data upon which that was based were so vague as to be useless.

Williams tasked John Opel, a future CEO of the company, to find out exactly what work-in-process inventory amounted to. Opel would give Williams an estimate only to find out within a day that he had not only missed, but missed by \$50 million. Opel finally got so frustrated that he insisted that each factory manager take a physical inventory, something IBM had never had to do before. He finally discovered that the accounting system had gone completely out of whack. IBM had almost \$600 million of work-in-process inventory that none of the factory managers wanted to claim. This made top management frantic.²⁴

Here is another example of how badly the forecasting was going. In 1963, IBM decided to pre-pay a major loan from the Prudential Insurance Company bearing an average interest rate of 3.5 percent. But when these unaccounted-for costs started popping up in such large amounts, it actually looked like IBM—the bastion of financial security—might run out of cash. In 1966, the company had to establish bank lines of credit for the same millions it had pre-paid earlier and had to pay about two percentage points more for any of the funds used. Truly startling was that, due to the cash shortage, IBM unexpectedly sold \$371 million of stock to the public in the spring of 1965.

The Price of Success

Tom Watson observed that in the autumn of 1965, “everything looked black, black, black. I was 51

years old, I had nine years of fantastic success behind me, and I didn't want my career to be wrecked by an announcement that the whole new product line was never going to fly at all. I panicked.”²⁵ Watson demoted his brother Dick and, because he believed that business needed a dictator to move forward, put the whole 360 program in the hands of Vin Learson. This decision, not surprisingly, alienated the brothers. Watson later reflected, “We remade the computer industry with the System/360, and objectively it was the greatest triumph of my business career. But whenever I think of it, I think about the brother I injured.”²⁶

We remade the computer industry with the System/360, and objectively it was the greatest triumph of my business career. But whenever I think of it, I think about the brother I injured.

THOMAS WATSON JR.

Here is Learson's assessment, “We made two miscalculations. We were off on our assessment of 360's potential reception, and we were off on our assessment of IBM's production capability to meet the demand. We did what Charles Kettering, an engineering genius and president of the General Motors Research Division, always advised against: we put a delivery date on something yet to be invented.”²⁷

IBM gambled \$5.25 billion on the 360, or 1.9 times the revenue for 1962. (When the 360 was approved, it was budgeted to cost \$675 million.)²⁸ The gamble turned out to be a fantastic success. But the price was high in terms not only of money, but of human relationships.

Perhaps too high. The company never succeeded at anything of this magnitude again. ○

²⁰ “Processor products : Final report of SPREAD task group,” IBM, December 28, 1961, Jerome Svigals papers, Lot X3951.2007, Box 5, Catalog Number 102713231, Computer History Museum.

²¹ Ingle and Bower, “Giant,” 7.

²² Thomas J. Watson Jr. and Peter Petre, *Father, Son & Co: My Life at IBM and Beyond* (New York: Bantam, 1990), 353.

²³ *Ibid.*, 350.

²⁴ *Ibid.*, 358.

²⁵ *Ibid.*, 357.

²⁶ *Ibid.*, 360.

²⁷ Rowena Olegario, “IBM and the Two Thomas J. Watsons” in Thomas K. McCraw, ed., *Creating Modern Capitalism* (Cambridge: Harvard University Press, 1997), 392.

²⁸ Boyer, *The 360 Revolution*, 31.