

---

# COLUMBIA

## ACCIDENT INVESTIGATION BOARD

---



---

REPORT VOLUME I  
AUGUST 2003

---

## AN INTRODUCTION TO NASA

“An Act to provide for research into the problems of flight within and outside the Earth’s atmosphere, and for other purposes.” With this simple preamble, the Congress and the President of the United States created the National Aeronautics and Space Administration (NASA) on October 1, 1958. Formed in response to the launch of *Sputnik* by the Soviet Union, NASA inherited the research-oriented National Advisory Committee for Aeronautics (NACA) and several other government organizations, and almost immediately began working on options for manned space flight. NASA’s first high profile program was Project Mercury, an early effort to learn if humans could survive in space. Project Gemini followed with a more complex series of experiments to increase man’s time in space and validate advanced concepts such as rendezvous. The efforts continued with Project Apollo, culminating in 1969 when *Apollo 11* landed the first humans on the Moon. The return from orbit on July 24, 1975, of the crew from the Apollo-Soyuz Test Project began a six-year hiatus of American manned space flight. The launch of the first Space Shuttle in April 1981 brought Americans back into space, continuing today with the assembly and initial operations of the International Space Station.

In addition to the human space flight program, NASA also maintains an active (if small) aeronautics research program, a space science program (including deep space and interplanetary exploration), and an Earth observation program. The agency also conducts basic research activities in a variety of fields.

NASA, like many federal agencies, is a heavily matrixed organization, meaning that the lines of authority are not necessarily straightforward. At the simplest level, there are three major types of entities involved in the Human Space Flight Program: NASA field centers, NASA programs carried out at those centers, and industrial and academic contractors. The centers provide the buildings, facilities, and support services for the various programs. The programs, along with field centers and Headquarters, hire civil servants and contractors from the private sector to support aspects of their enterprises.

### THE LOCATIONS

NASA Headquarters, located in Washington D.C., is responsible for leadership and management across five strategic enterprises: Aerospace Technology, Biological and Physical Research, Earth Science, Space Science, and Human Exploration and Development of Space. NASA Headquarters also provides strategic management for the Space Shuttle and International Space Station programs.

The Johnson Space Center in Houston, Texas, was established in 1961 as the Manned Spacecraft Center and has led the development of every U.S. manned space flight program. Currently, Johnson is home to both the Space Shuttle and International Space Station Program Offices. The facilities at Johnson include the training, simulation, and mission control centers for the Space Shuttle and Space Station. Johnson also has flight operations at Ellington Field, where the training aircraft for the astronauts and support aircraft for the Space Shuttle Program are stationed, and manages the White Sands Test Facility, New Mexico, where hazardous testing is conducted.

The Kennedy Space Center was created to launch the Apollo missions to the Moon, and currently provides launch and landing facilities for the Space Shuttle. The Center is located on Merritt Island, Florida, adjacent to the Cape Canaveral Air Force Station that also provides support for the Space Shuttle Program (and was the site of the earlier Mercury and Gemini launches). Personnel at Kennedy support maintenance and overhaul services for the Orbiters, assemble and check-out the integrated vehicle prior to launch, and operate the Space Station Processing Facility where components of the orbiting laboratory are packaged for launch aboard the Space Shuttle. The majority of contractor personnel assigned to Kennedy are part of the Space Flight Operations Contract administered by the Space Shuttle Program Office at Johnson.

The Marshall Space Flight Center, near Huntsville, Alabama, is home to most NASA rocket propulsion efforts. The Space Shuttle Projects Office located at Marshall—organizationally part of the Space Shuttle Program Office at Johnson—manages the manufacturing and support contracts to Boeing Rocketdyne for the Space Shuttle Main Engine (SSME), to Lockheed Martin for the External Tank (ET), and to ATK Thiokol Propulsion for the Reusable Solid Rocket Motor (RSRM, the major piece of the Solid Rocket Booster). Marshall is also involved in micro-gravity research and space product development programs that fly as payloads on the Space Shuttle.



The Stennis Space Center in Bay St. Louis, Mississippi, is the largest rocket propulsion test complex in the United States. Stennis provides all of the testing facilities for the Space



# The Evolution of the Space Shuttle Program

More than two decades after its first flight, the Space Shuttle remains the only reusable spacecraft in the world capable of simultaneously putting multiple-person crews and heavy cargo into orbit, of deploying, servicing, and retrieving satellites, and of returning the products of on-orbit research to Earth. These capabilities are an important asset for the United States and its international partners in space. Current plans call for the Space Shuttle to play a central role in the U.S. human space flight program for years to come.

The Space Shuttle Program's remarkable successes, however, come with high costs and tremendous risks. The February 1 disintegration of *Columbia* during re-entry, 17 years after *Challenger* was destroyed on ascent, is the most recent reminder that sending people into orbit and returning them safely to Earth remains a difficult and perilous endeavor.

It is the view of the Columbia Accident Investigation Board that the *Columbia* accident is not a random event, but rather a product of the Space Shuttle Program's history and current management processes. Fully understanding how it happened requires an exploration of that history and management. This chapter charts how the Shuttle emerged from a series of political compromises that produced unreasonable expectations – even myths – about its performance, how the *Challenger* accident shattered those myths several years after NASA began acting upon them as fact, and how, in retrospect, the Shuttle's technically ambitious design resulted in an inherently vulnerable vehicle, the safe operation of which exceeded NASA's organizational capabilities as they existed at the time of the *Columbia* accident. The Board's investigation of what caused the *Columbia* accident thus begins in the fields of East Texas but reaches more than 30 years into the past, to a series of economically and politically driven decisions that cast the Shuttle program in a role that its nascent technology could not support. To understand the cause of the *Columbia* accident is to understand how a program promising reliability and cost efficiency resulted instead in a developmental vehicle that never achieved the fully operational status NASA and the nation accorded it.

## 1.1 GENESIS OF THE SPACE TRANSPORTATION SYSTEM

The origins of the Space Shuttle Program date to discussions on what should follow Project Apollo, the dramatic U.S. missions to the moon.<sup>1</sup> NASA centered its post-Apollo plans on developing increasingly larger outposts in Earth orbit that would be launched atop Apollo's immense Saturn V booster. The space agency hoped to construct a 12-person space station by 1975; subsequent stations would support 50, then 100 people. Other stations would be placed in orbit around the moon and then be constructed on the lunar surface. In parallel, NASA would develop the capability for the manned exploration of Mars. The concept of a vehicle – or Space Shuttle – to take crews and supplies to and from low-Earth orbit arose as part of this grand vision (see Figure 1.1-1). To keep the costs of these trips to a minimum, NASA intended to develop a fully reusable vehicle.<sup>2</sup>



Figure 1.1-1. Early concepts for the Space Shuttle envisioned a reusable two-stage vehicle with the reliability and versatility of a commercial airliner.

hicle has proved difficult and costly to operate, riskier than expected, and, on two occasions, deadly.

It is the Board's view that, in retrospect, the increased complexity of a Shuttle designed to be all things to all people created inherently greater risks than if more realistic technical goals had been set at the start. Designing a reusable spacecraft that is also cost-effective is a daunting engineering challenge; doing so on a tightly constrained budget is even more difficult. Nevertheless, the remarkable system we have today is a reflection of the tremendous engineering expertise and dedication of the workforce that designed and built the Space Shuttle within the constraints it was given.

In the end, the greatest compromise NASA made was not so much with any particular element of the technical design, but rather with the premise of the vehicle itself. NASA promised it could develop a Shuttle that would be launched almost on demand and would fly many missions each year. Throughout the history of the program, a gap has persisted between the rhetoric NASA has used to market the Space Shuttle and operational reality, leading to an enduring image of the Shuttle as capable of safely and routinely carrying out missions with little risk.

### 1.3 SHUTTLE DEVELOPMENT, TESTING, AND QUALIFICATION

The Space Shuttle was subjected to a variety of tests before its first flight. However, NASA conducted these tests somewhat differently than it had for previous spacecraft.<sup>8</sup> The Space Shuttle Program philosophy was to ground-test key hardware elements such as the main engines, Solid Rocket Boosters, External Tank, and Orbiter separately and to use analytical models, not flight testing, to certify the integrated Space Shuttle system. During the Approach and Landing Tests (see Figure 1.3-1), crews verified that the Orbiter could successfully fly at low speeds and land safely; however, the Space Shuttle was not flown on an unmanned orbital test flight prior to its first mission – a significant change in philosophy compared to that of earlier American spacecraft.



Figure 1.3-1. The first Orbiter was Enterprise, shown here being released from the Boeing 747 Shuttle Carrier Aircraft during the Approach and Landing Tests at Edwards Air Force Base.

The significant advances in technology that the Shuttle's design depended on led its development to run behind schedule. The date for the first Space Shuttle launch slipped from March 1978 to 1979, then to 1980, and finally to the spring of 1981. One historian has attributed one year of this delay "to budget cuts, a second year to problems with the main engines, and a third year to problems with the thermal protection tiles."<sup>9</sup> Because of these difficulties, in 1979 the program underwent an exhaustive White House review. The program was thought to be a billion dollars over budget, and President Jimmy Carter wanted to make sure that it was worth continuing. A key factor in the White House's final assessment was that the Shuttle was needed to launch the intelligence satellites required for verification of the SALT II arms control treaty, a top Carter Administration priority. The review reaffirmed the need for the Space Shuttle, and with continued White House and Congressional support, the path was clear for its transition from development to flight. NASA ultimately completed Shuttle development for only 15 percent more than its projected cost, a comparatively small cost overrun for so complex a program.<sup>10</sup>

The Orbiter that was destined to be the first to fly into space was *Columbia*. In early 1979, NASA was beginning to feel the pressure of being behind schedule. Despite the fact that only 24,000 of the 30,000 Thermal Protection System tiles had been installed, NASA decided to fly *Columbia* from the manufacturing plant in Palmdale, California, to the Kennedy Space Center in March 1979. The rest of the tiles would be installed in Florida, thus allowing NASA to maintain the appearance of *Columbia's* scheduled launch date. Problems with the main engines and the tiles were to leave *Columbia* grounded for two more years.

### 1.4 THE SHUTTLE BECOMES "OPERATIONAL"

On the first Space Shuttle mission, STS-1,<sup>11</sup> *Columbia* carried John W. Young and Robert L. Crippen to orbit on April 12, 1981, and returned them safely two days later to Edwards Air Force Base in California (see Figure 1.4-1). After three years of policy debate and nine years of development, the Shuttle returned U.S. astronauts to space for the first time since the Apollo-Soyuz Test Project flew in July 1975. Post-flight inspection showed that *Columbia* suffered slight damage from excess Solid Rocket Booster ignition pressure and lost 16 tiles, with 148 others sustaining some damage. Over the following 15 months, *Columbia* was launched three more times. At the end of its fourth mission, on July 4, 1982, *Columbia* landed at Edwards where President Ronald Reagan declared to a nation celebrating Independence Day that "beginning with the next flight, the *Columbia* and her sister ships will be *fully operational*, ready to provide *economical and routine access to space* for scientific exploration, commercial ventures, and for tasks related to the national security" [emphasis added].<sup>12</sup>

There were two reasons for declaring the Space Shuttle "operational" so early in its flight program. One was NASA's hope for quick Presidential approval of its next manned space flight program, a space station, which would not move forward while the Shuttle was still considered developmental. The second reason was that the nation was sud-



# Accident Analysis

One of the central purposes of this investigation, like those for other kinds of accidents, was to identify the chain of circumstances that caused the *Columbia* accident. In this case the task was particularly challenging, because the breakup of the Orbiter occurred at hypersonic velocities and extremely high altitudes, and the debris was scattered over a wide area. Moreover, the initiating event preceded the accident by more than two weeks. In pursuit of the sequence of the cause, investigators developed a broad array of information sources. Evidence was derived from film and video of the launch, radar images of *Columbia* on orbit, and amateur video of debris shedding during the in-flight breakup. Data was obtained from sensors onboard the Orbiter – some of this data was downlinked during the flight, and some came from an on-board recorder that was recovered during the debris search. Analysis of the debris was particularly valuable to the investigation. Clues were to be found not only in the condition of the pieces, but also in their location – both where they had been on the Orbiter and where they were found on the ground. The investigation also included extensive computer modeling, impact tests, wind tunnel studies, and other analytical techniques. Each of these avenues of inquiry is described in this chapter.

Because it became evident that the key event in the chain leading to the accident involved both the External Tank and one of the Orbiter's wings, the chapter includes a study of these two structures. The understanding of the accident's physical cause that emerged from this investigation is summarized in the statement at the beginning of the chapter. Included in the chapter are the findings and recommendations of the Columbia Accident Investigation Board that are based on this examination of the physical evidence.

## 3.1 THE PHYSICAL CAUSE

The physical cause of the loss of *Columbia* and its crew was a breach in the Thermal Protection System on the leading edge of the left wing. The breach was initiated by a piece of insulating foam that separated from the left bipod ramp of the External Tank and struck the wing in the vicinity of the lower half of Reinforced Carbon-Carbon panel 8 at 81.9 seconds after launch. During re-entry, this breach in the Thermal

Protection System allowed superheated air to penetrate the leading-edge insulation and progressively melt the aluminum structure of the left wing, resulting in a weakening of the structure until increasing aerodynamic forces caused loss of control, failure of the wing, and breakup of the Orbiter.



Figure 3.1-1. Columbia sitting at Launch Complex 39-A. The upper circle shows the left bipod (-Y) ramp on the forward attach point, while the lower circle is around RCC panel 8-left.



# Other Factors Considered

During its investigation, the Board evaluated every known factor that could have caused or contributed to the *Columbia* accident, such as the effects of space weather on the Orbiter during re-entry and the specters of sabotage and terrorism. In addition to the analysis/scenario investigations, the Board oversaw a NASA “fault tree” investigation, which accounts for every chain of events that could possibly cause a system to fail. Most of these factors were conclusively eliminated as having nothing to do with the accident; however, several factors have yet to be ruled out. Although deemed by the Board as unlikely to have contributed to the accident, these are still open and are being investigated further by NASA. In a few other cases, there is insufficient evidence to completely eliminate a factor, though most evidence indicates that it did not play a role in the accident. In the course of investigating these factors, the Board identified several serious problems that were not part of the accident’s causal chain but nonetheless have major implications for future missions.

In this chapter, a discussion of these potential causal and contributing factors is divided into two sections. The first introduces the primary tool used to assess potential causes of the breakup: the fault tree. The second addresses fault tree items and particularly notable factors that raised concerns for this investigation and, more broadly, for the future operation of the Space Shuttle.

## 4.1 FAULT TREE

The NASA Accident Investigation Team investigated the accident using “fault trees,” a common organizational tool in systems engineering. Fault trees are graphical representations of every conceivable sequence of events that could cause a system to fail. The fault tree’s uppermost level illustrates the events that could have directly caused the loss of *Columbia* by aerodynamic breakup during re-entry. Subsequent levels comprise all individual elements or factors that could cause the failure described immediately above it. In this way, all potential chains of causation that lead ultimately to the loss of *Columbia* can be diagrammed, and the behavior of every subsystem that was not a precipitating cause can be eliminated from consideration. Figure 4.1-1 depicts the fault tree structure for the *Columbia* accident investigation.

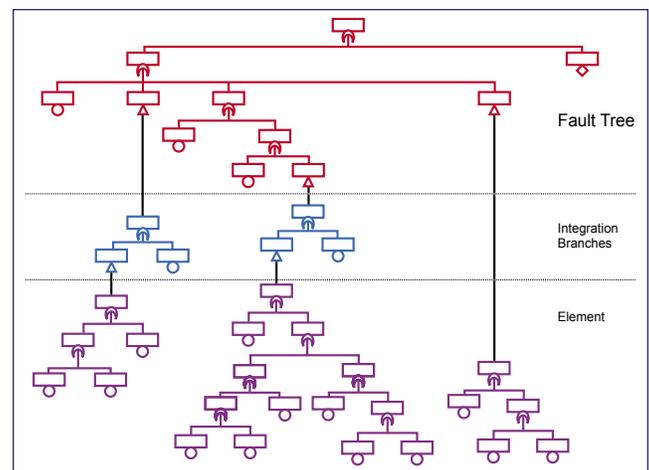


Figure 4.1-1. Accident investigation fault tree structure.

NASA chartered six teams to develop fault trees, one for each of the Shuttle’s major components: the Orbiter, Space Shuttle Main Engine, Reusable Solid Rocket Motor, Solid Rocket Booster, External Tank, and Payload. A seventh “systems integration” fault tree team analyzed failure scenarios involving two or more Shuttle components. These interdisciplinary teams included NASA and contractor personnel, as well as outside experts.

Some of the fault trees are very large and intricate. For instance, the Orbiter fault tree, which only considers events on the Orbiter that could have led to the accident, includes 234 elements. In contrast, the Systems Integration fault tree, which deals with interactions among parts of the Shuttle, includes 295 unique multi-element integration faults, 128 Orbiter multi-element faults, and 221 connections to the other Shuttle components. These faults fall into three categories: induced and natural environments (such as structural interface loads and electromechanical effects); integrated vehicle mass properties; and external impacts (such as debris from the External Tank). Because the Systems Integration team considered multi-element faults – that is, scenarios involving several Shuttle components – it frequently worked in tandem with the Component teams.



# Part Two

## Why The Accident Occurred

Many accident investigations do not go far enough. They identify the technical cause of the accident, and then connect it to a variant of “operator error” – the line worker who forgot to insert the bolt, the engineer who miscalculated the stress, or the manager who made the wrong decision. But this is seldom the entire issue. When the determinations of the causal chain are limited to the technical flaw and individual failure, typically the actions taken to prevent a similar event in the future are also limited: fix the technical problem and replace or retrain the individual responsible. Putting these corrections in place leads to another mistake – the belief that the problem is solved. The Board did not want to make these errors.

Attempting to manage high-risk technologies while minimizing failures is an extraordinary challenge. By their nature, these complex technologies are intricate, with many interrelated parts. Standing alone, the components may be well understood and have failure modes that can be anticipated. Yet when these components are integrated into a larger system, unanticipated interactions can occur that lead to catastrophic outcomes. The risk of these complex systems is increased when they are produced and operated by complex organizations that also break down in unanticipated ways.

In our view, the NASA organizational culture had as much to do with this accident as the foam. Organizational culture refers to the basic values, norms, beliefs, and practices that characterize the functioning of an institution. At the most basic level, organizational culture defines the assumptions that employees make as they carry out their work. It is a powerful force that can persist through reorganizations and the change of key personnel. It can be a positive or a negative force.

In a report dealing with nuclear wastes, the National Research Council quoted Alvin Weinberg’s classic statement about the “Faustian bargain” that nuclear scientists made with society. “The price that we demand of society for this magical energy source is both a vigilance and a longevity of our social institutions that we are quite unaccustomed to.” This is also true of the space program. At NASA’s urging, the nation committed to building an amazing, if compromised,

vehicle called the Space Shuttle. When the agency did this, it accepted the bargain to operate and maintain the vehicle in the safest possible way. The Board is not convinced that NASA has completely lived up to the bargain, or that Congress and the Administration has provided the funding and support necessary for NASA to do so. This situation needs to be addressed – if the nation intends to keep conducting human space flight, it needs to live up to its part of the bargain.

Part Two of this report examines NASA’s organizational, historical, and cultural factors, as well as how these factors contributed to the accident. As in Part One, this part begins with history. Chapter 5 examines the post-*Challenger* history of NASA and its Human Space Flight Program. This includes reviewing the budget as well as organizational and management history, such as shifting management systems and locations. Chapter 6 documents management performance related to *Columbia* to establish events analyzed in later chapters. The chapter reviews the foam strikes, intense schedule pressure driven by an artificial requirement to deliver Node 2 to the International Space Station by a certain date, and NASA management’s handling of concerns regarding *Columbia* during the STS-107 mission.

In Chapter 7, the Board presents its views of how high-risk activities should be managed, and lists the characteristics of institutions that emphasize high-reliability results over economic efficiency or strict adherence to a schedule. This chapter measures the Space Shuttle Program’s organizational and management practices against these principles and finds them wanting. Chapter 7 defines the organizational cause and offers recommendations. Chapter 8 draws from the previous chapters on history, budgets, culture, organization, and safety practices, and analyzes how all these factors contributed to this accident. This chapter captures the Board’s views of the need to adjust management to enhance safety margins in Shuttle operations, and reaffirms the Board’s position that without these changes, we have no confidence that other “corrective actions” will improve the safety of Shuttle operations. The changes we recommend will be difficult to accomplish – and will be internally resisted.



# From Challenger to Columbia

The Board is convinced that the factors that led to the *Columbia* accident go well beyond the physical mechanisms discussed in Chapter 3. The causal roots of the accident can also be traced, in part, to the turbulent post-Cold War policy environment in which NASA functioned during most of the years between the destruction of *Challenger* and the loss of *Columbia*. The end of the Cold War in the late 1980s meant that the most important political underpinning of NASA's Human Space Flight Program – U.S.-Soviet space competition – was lost, with no equally strong political objective to replace it. No longer able to justify its projects with the kind of urgency that the superpower struggle had provided, the agency could not obtain budget increases through the 1990s. Rather than adjust its ambitions to this new state of affairs, NASA continued to push an ambitious agenda of space science and exploration, including a costly Space Station Program.

If NASA wanted to carry out that agenda, its only recourse, given its budget allocation, was to become more efficient, accomplishing more at less cost. The search for cost reductions led top NASA leaders over the past decade to downsize the Shuttle workforce, outsource various Shuttle Program responsibilities – including safety oversight – and consider eventual privatization of the Space Shuttle Program. The program's budget was reduced by 40 percent in purchasing power over the past decade and repeatedly raided to make up for Space Station cost overruns, even as the Program maintained a launch schedule in which the Shuttle, a developmental vehicle, was used in an operational mode. In addition, the uncertainty of top policymakers in the White House, Congress, and NASA as to how long the Shuttle would fly before being replaced resulted in the delay of upgrades needed to make the Shuttle safer and to extend its service life.

The Space Shuttle Program has been transformed since the late 1980s implementation of post-*Challenger* management changes in ways that raise questions, addressed here and in later chapters of Part Two, about NASA's ability to safely

operate the Space Shuttle. While it would be inaccurate to say that NASA managed the Space Shuttle Program at the time of the *Columbia* accident in the same manner it did prior to *Challenger*, there are unfortunate similarities between the agency's performance and safety practices in both periods.

## 5.1 THE CHALLENGER ACCIDENT AND ITS AFTERMATH

The inherently vulnerable design of the Space Shuttle, described in Chapter 1, was a product of policy and technological compromises made at the time of its approval in 1972. That approval process also produced unreasonable expectations, even myths, about the Shuttle's future performance that NASA tried futilely to fulfill as the Shuttle became "operational" in 1982. At first, NASA was able to maintain the image of the Shuttle as an operational vehicle. During its early years of operation, the Shuttle launched satellites, performed on-orbit research, and even took members of Congress into orbit. At the beginning of 1986, the goal of "routine access to space" established by President Ronald Reagan in 1982 was ostensibly being achieved. That appearance soon proved illusory. On the cold morning of January 28, 1986, the Shuttle *Challenger* broke apart 73 seconds into its climb towards orbit. On board were Francis R. Scobee, Michael J. Smith, Ellison S. Onizuka, Judith A. Resnick, Ronald E. McNair, Sharon Christa McAuliffe, and Gregory B. Jarvis. All perished.

### Rogers Commission

On February 3, 1986, President Reagan created the Presidential Commission on the Space Shuttle Challenger Accident, which soon became known as the Rogers Commission after its chairman, former Secretary of State William Rogers. The Commission's report, issued on June 6, 1986, concluded that the loss of *Challenger* was caused by a failure of the joint and seal between the two lower segments of the right Solid Rocket Booster. Hot gases blew past a rubber O-ring in the joint, leading to a structural failure and the explosive burn-

discussed in this chapter, there were at least three major contributing factors to that environment:

- Throughout the decade, the Shuttle Program has had to function within an increasingly constrained budget. Both the Shuttle budget and workforce have been reduced by over 40 percent during the past decade. The White House, Congress, and NASA leadership exerted constant pressure to reduce or at least freeze operating costs. As a result, there was little margin in the budget to deal with unexpected technical problems or make Shuttle improvements.
- The Shuttle was mischaracterized by the 1995 Kraft Report as “a mature and reliable system ... about as safe as today’s technology will provide.” Based on this mischaracterization, NASA believed that it could turn increased responsibilities for Shuttle operations over to a single prime contractor and reduce its direct involvement in ensuring safe Shuttle operations, instead monitoring contractor performance from a more detached position. NASA also believed that it could use the “mature” Shuttle to carry out operational missions without continually focusing engineering attention on understanding the mission-by-mission anomalies inherent in a developmental vehicle.
- In the 1990s, the planned date for replacing the Shuttle shifted from 2006 to 2012 and then to 2015 or later. Given the uncertainty regarding the Shuttle’s service life, there has been policy and budgetary ambivalence on investing in the vehicle. Only in the past year has NASA begun to provide the resources needed to sustain extended Shuttle operations. Previously, safety and support upgrades were delayed or deferred, and Shuttle infrastructure was allowed to deteriorate.

The Board observes that this is hardly an environment in which those responsible for safe operation of the Shuttle can function without being influenced by external pressures. It is to the credit of Space Shuttle managers and the Shuttle workforce that the vehicle was able to achieve its program objectives for as long as it did.

An examination of the Shuttle Program’s history from *Challenger* to *Columbia* raises the question: Did the Space Shuttle Program budgets constrained by the White House and Congress threaten safe Shuttle operations? There is no straightforward answer. In 1994, an analysis of the Shuttle budget concluded that reductions made in the early 1990s represented a “healthy tightening up” of the program.<sup>77</sup> Certainly those in the Office of Management and Budget and in NASA’s congressional authorization and appropriations subcommittees thought they were providing enough resources to operate the Shuttle safely, while also taking into account the expected Shuttle lifetime and the many other demands on the Federal budget. NASA Headquarters agreed, at least until Administrator Goldin declared a “space launch crisis” in June 1999 and asked that additional resources for safety upgrades be added to the NASA budget. By 2001, however, one experienced observer of the space program described the Shuttle workforce as “The Few, the Tired,”

and suggested that “a decade of downsizing and budget tightening has left NASA exploring the universe with a less experienced staff and older equipment.”<sup>78</sup>

It is the Board’s view that this latter statement is an accurate depiction of the Space Shuttle Program at the time of STS-107. The Program was operating too close to too many margins. The Board also finds that recent modest increases in the Shuttle Program’s budget are necessary and overdue steps toward providing the resources to sustain the program for its now-extended lifetime. Similarly, NASA has recently recognized that providing an adequately sized and appropriately trained workforce is critical to the agency’s future success.

An examination of the Program’s management changes also leads to the question: Did turmoil in the management structure contribute to the accident? The Board found no evidence that the transition from many Space Shuttle contractors to a partial consolidation of contracts under a single firm has by itself introduced additional technical risk into the Space Shuttle Program. The transfer of responsibilities that has accompanied the Space Flight Operations Contract has, however, complicated an already complex Program structure and created barriers to effective communication. Designating the Johnson Space Center as the “lead center” for the Space Shuttle Program did resurrect some of the Center rivalries and communication difficulties that existed before the *Challenger* accident. The specific ways in which this complexity and lack of an integrated approach to Shuttle management impinged on NASA’s performance during and before the flight of STS-107 are discussed in Chapters 6 and 7.

As the 21st century began, NASA’s deeply ingrained human space flight culture – one that has evolved over 30 years as the basis for a more conservative, less technically and organizationally capable organization than the Apollo-era NASA – remained strong enough to resist external pressures for adaptation and change. At the time of the launch of STS-107, NASA retained too many negative (and also many positive) aspects of its traditional culture: “flawed decision making, self deception, introversion and a diminished curiosity about the world outside the perfect place.”<sup>79</sup> These characteristics were reflected in NASA’s less than stellar performance before and during the STS-107 mission, which is described in the following chapters.



# Decision Making at NASA

The dwindling post-Cold War Shuttle budget that launched NASA leadership on a crusade for efficiency in the decade before *Columbia*'s final flight powerfully shaped the environment in which Shuttle managers worked. The increased organizational complexity, transitioning authority structures, and ambiguous working relationships that defined the restructured Space Shuttle Program in the 1990s created turbulence that repeatedly influenced decisions made before and during STS-107.

This chapter connects Chapter 5's analysis of NASA's broader policy environment to a focused scrutiny of Space Shuttle Program decisions that led to the STS-107 accident. Section 6.1 illustrates how foam debris losses that violated design requirements came to be defined by NASA management as an acceptable aspect of Shuttle missions, one that posed merely a maintenance "turnaround" problem rather than a safety-of-flight concern. Section 6.2 shows how, at a pivotal juncture just months before the *Columbia* accident, the management goal of completing Node 2 of the International Space Station on time encouraged Shuttle managers to continue flying, even after a significant bipod-foam debris strike on STS-112. Section 6.3 notes the decisions made during STS-107 in response to the bipod foam strike, and reveals how engineers' concerns about risk and safety were competing with – and were defeated by – management's belief that foam could not hurt the Orbiter, as well as the need to keep on schedule. In relating a rescue and repair scenario that might have enabled the crew's safe return, Section 6.4 grapples with yet another latent assumption held by Shuttle managers during and after STS-107: that even if the foam strike had been discovered, nothing could have been done.

## 6.1 A HISTORY OF FOAM ANOMALIES

The shedding of External Tank foam – the physical cause of the *Columbia* accident – had a long history. Damage caused by debris has occurred on every Space Shuttle flight, and most missions have had insulating foam shed during ascent. This raises an obvious question: Why did NASA continue

flying the Shuttle with a known problem that violated design requirements? It would seem that the longer the Shuttle Program allowed debris to continue striking the Orbiters, the more opportunity existed to detect the serious threat it posed. But this is not what happened. Although engineers have made numerous changes in foam design and application in the 25 years that the External Tank has been in production, the problem of foam-shedding has not been solved, nor has the Orbiter's ability to tolerate impacts from foam or other debris been significantly improved.

### The Need for Foam Insulation

The External Tank contains liquid oxygen and hydrogen propellants stored at minus 297 and minus 423 degrees Fahrenheit. Were the super-cold External Tank not sufficiently insulated from the warm air, its liquid propellants would boil, and atmospheric nitrogen and water vapor would condense and form thick layers of ice on its surface. Upon launch, the ice could break off and damage the Orbiter. (See Chapter 3.)

To prevent this from happening, large areas of the External Tank are machine-sprayed with one or two inches of foam, while specific fixtures, such as the bipod ramps, are hand-sculpted with thicker coats. Most of these insulating materials fall into a general category of "foam," and are outwardly similar to hardware store-sprayable foam insulation. The problem is that foam does not always stay where the External Tank manufacturer Lockheed Martin installs it. During flight, popcorn- to briefcase-size chunks detach from the External Tank.

### Original Design Requirements

Early in the Space Shuttle Program, foam loss was considered a dangerous problem. Design engineers were extremely concerned about potential damage to the Orbiter and its fragile Thermal Protection System, parts of which are so vulnerable to impacts that lightly pressing a thumbnail into them leaves a mark. Because of these concerns, the baseline



# The Accident's Organizational Causes

Many accident investigations make the same mistake in defining causes. They identify the widget that broke or malfunctioned, then locate the person most closely connected with the technical failure: the engineer who miscalculated an analysis, the operator who missed signals or pulled the wrong switches, the supervisor who failed to listen, or the manager who made bad decisions. When causal chains are limited to technical flaws and individual failures, the ensuing responses aimed at preventing a similar event in the future are equally limited: they aim to fix the technical problem and replace or retrain the individual responsible. Such corrections lead to a misguided and potentially disastrous belief that the underlying problem has been solved. The Board did not want to make these errors. A central piece of our expanded cause model involves NASA as an organizational whole.

## ORGANIZATIONAL CAUSE STATEMENT

The organizational causes of this accident are rooted in the Space Shuttle Program's history and culture, including the original compromises that were required to gain approval for the Shuttle Program, subsequent years of resource constraints, fluctuating priorities, schedule pressures, mischaracterizations of the Shuttle as operational rather than developmental, and lack of an agreed national vision. Cultural traits and organizational practices detrimental to safety and reliability were allowed to develop, including: reliance on past success as a substitute for sound engineering practices (such as testing to understand why systems were not performing in accordance with requirements/specifications); organizational barriers which prevented effective communication of critical safety information and stifled professional differences of opinion; lack of integrated management across program elements; and the evolution of an informal chain of command and decision-making processes that operated outside the organization's rules.

## UNDERSTANDING CAUSES

In the Board's view, NASA's organizational culture and structure had as much to do with this accident as the External Tank foam. Organizational culture refers to the values, norms, beliefs, and practices that govern how an institution functions. At the most basic level, organizational culture defines the assumptions that employees make as they carry out their work. It is a powerful force that can persist through reorganizations and the reassignment of key personnel.

Given that today's risks in human space flight are as high and the safety margins as razor thin as they have ever been, there is little room for overconfidence. Yet the attitudes and decision-making of Shuttle Program managers and engineers during the events leading up to this accident were clearly overconfident and often bureaucratic in nature. They deferred to layered and cumbersome regulations rather than the fundamentals of safety. The Shuttle Program's safety culture is straining to hold together the vestiges of a once robust systems safety program.

As the Board investigated the *Columbia* accident, it expected to find a vigorous safety organization, process, and culture at NASA, bearing little resemblance to what the Rogers Commission identified as the ineffective "silent safety" system in which budget cuts resulted in a lack of resources, personnel, independence, and authority. NASA's initial briefings to the Board on its safety programs espoused a risk-averse philosophy that empowered any employee to stop an operation at the mere glimmer of a problem. Unfortunately, NASA's views of its safety culture in those briefings did not reflect reality. Shuttle Program safety personnel failed to adequately assess anomalies and frequently accepted critical risks without qualitative or quantitative support, even when the tools to provide more comprehensive assessments were available.

Similarly, the Board expected to find NASA's Safety and Mission Assurance organization deeply engaged at every

petitive sourcing options for the Shuttle Program. In its final report to NASA, the team highlighted several safety-related concerns, which the Board shares:

- Flight and ground hardware and software are obsolete, and safety upgrades and aging infrastructure repairs have been deferred.
- Budget constraints have impacted personnel and resources required for maintenance and upgrades.
- International Space Station schedules exert significant pressures on the Shuttle Program.
- Certain mechanisms may impede worker anonymity in reporting safety concerns.
- NASA does not have a truly independent safety function with the authority to halt the progress of a critical mission element.<sup>11</sup>

Based on these findings, the task force suggested that an Independent Safety Assurance function should be created that would hold one of “three keys” in the Certification of Flight Readiness process (NASA and the operating contractor would hold the other two), effectively giving this function the ability to stop any launch. Although in the Board’s view the “third key” Certification of Flight Readiness process is not a perfect solution, independent safety and verification functions are vital to continued Shuttle operations. This independent function should possess the authority to shut down the flight preparation processes or intervene post-launch when an anomaly occurs.

## 7.2 ORGANIZATIONAL CAUSES: INSIGHTS FROM THEORY

To develop a thorough understanding of accident causes and risk, and to better interpret the chain of events that led to the *Columbia* accident, the Board turned to the contemporary social science literature on accidents and risk and sought insight from experts in High Reliability, Normal Accident, and Organizational Theory.<sup>12</sup> Additionally, the Board held a forum, organized by the National Safety Council, to define the essential characteristics of a sound safety program.<sup>13</sup>

High Reliability Theory argues that organizations operating high-risk technologies, if properly designed and managed, can compensate for inevitable human shortcomings, and therefore avoid mistakes that under other circumstances would lead to catastrophic failures.<sup>14</sup> Normal Accident Theory, on the other hand, has a more pessimistic view of the ability of organizations and their members to manage high-risk technology. Normal Accident Theory holds that organizational and technological complexity contributes to failures. Organizations that aspire to failure-free performance are inevitably doomed to fail because of the inherent risks in the technology they operate.<sup>15</sup> Normal Accident models also emphasize systems approaches and systems thinking, while the High Reliability model works from the bottom up: if each component is highly reliable, then the system will be highly reliable and safe.

Though neither High Reliability Theory nor Normal Accident Theory is entirely appropriate for understanding this accident, insights from each figured prominently in the

Board’s deliberation. Fundamental to each theory is the importance of strong organizational culture and commitment to building successful safety strategies.

The Board selected certain well-known traits from these models to use as a yardstick to assess the Space Shuttle Program, and found them particularly useful in shaping its views on whether NASA’s current organization of its Human Space Flight Program is appropriate for the remaining years of Shuttle operation and beyond. Additionally, organizational theory, which encompasses organizational culture, structure, history, and hierarchy, is used to explain the *Columbia* accident, and, ultimately, combines with Chapters 5 and 6 to produce an expanded explanation of the accident’s causes.<sup>16</sup> The Board believes the following considerations are critical to understand what went wrong during STS-107. They will become the central motifs of the Board’s analysis later in this chapter.

- **Commitment to a Safety Culture:** NASA’s safety culture has become reactive, complacent, and dominated by unjustified optimism. Over time, slowly and unintentionally, independent checks and balances intended to increase safety have been eroded in favor of detailed processes that produce massive amounts of data and unwarranted consensus, but little effective communication. Organizations that successfully deal with high-risk technologies create and sustain a disciplined safety system capable of identifying, analyzing, and controlling hazards throughout a technology’s life cycle.
- **Ability to Operate in Both a Centralized and Decentralized Manner:** The ability to operate in a centralized manner when appropriate, and to operate in a decentralized manner when appropriate, is the hallmark of a high-reliability organization. On the operational side, the Space Shuttle Program has a highly centralized structure. Launch commit criteria and flight rules govern every imaginable contingency. The Mission Control Center and the Mission Management Team have very capable decentralized processes to solve problems that are not covered by such rules. The process is so highly regarded that it is considered one of the best problem-solving organizations of its type.<sup>17</sup> In these situations, mature processes anchor rules, procedures, and routines to make the Shuttle Program’s matrixed workforce seamless, at least on the surface.

Nevertheless, it is evident that the position one occupies in this structure makes a difference. When supporting organizations try to “push back” against centralized Program direction – like the Debris Assessment Team did during STS-107 – independent analysis generated by a decentralized decision-making process can be stifled. The Debris Assessment Team, working in an essentially decentralized format, was well-led and had the right expertise to work the problem, but their charter was “fuzzy,” and the team had little direct connection to the Mission Management Team. This lack of connection to the Mission Management Team and the Mission Evaluation Room is the single most compelling reason why communications were so poor during the debris

## ENGINEERING BY VIEWGRAPHS

The Debris Assessment Team presented its analysis in a formal briefing to the Mission Evaluation Room that relied on PowerPoint slides from Boeing. When engineering analyses and risk assessments are condensed to fit on a standard form or overhead slide, information is inevitably lost. In the process, the priority assigned to information can be easily misrepresented by its placement on a chart and the language that is used. Dr. Edward Tufte of Yale University, an expert in information presentation who also researched communications failures in the *Challenger* accident, studied how the slides used by the Debris Assessment Team in their briefing to the Mission Evaluation Room misrepresented key information.<sup>38</sup>

The slide created six levels of hierarchy, signified by the title and the symbols to the left of each line. These levels prioritized information that was already contained in 11 simple sentences. Tufte also notes that the title is confusing. "Review of Test Data Indicates Conservatism" refers not to the predicted tile damage, but to the choice of test models used to predict the damage.

Only at the bottom of the slide do engineers state a key piece of information: that one estimate of the debris that struck *Columbia* was 640 times larger than the data used to calibrate the model on which engineers based their damage assessments. (Later analysis showed that the debris object was actually 400 times larger). This difference led Tufte to suggest that a more appropriate headline would be "Review of Test Data Indicates Irrelevance of Two Models."<sup>39</sup>

Tufte also criticized the sloppy language on the slide. "The vaguely quantitative words 'significant' and 'significantly' are used 5 times on this slide," he notes, "with de facto meanings ranging from 'detectable in largely irrelevant calibration case study' to 'an amount of damage so that everyone dies' to 'a difference of 640-fold.'" <sup>40</sup> Another example of sloppiness is that "cubic inches" is written inconsistently: "3cu. In.," "1920cu in.," and "3 cu in." While such inconsistencies might seem minor, in highly technical fields like aerospace engineering a misplaced decimal point or mistaken unit of measurement can easily engender inconsistencies and inaccuracies. In another phrase "Test results do show that it is possible at sufficient mass and velocity," the word "it" actually refers to "damage to the protective tiles."

As information gets passed up an organization hierarchy, from people who do analysis to mid-level managers to high-level leadership, key explanations and supporting information is filtered out. In this context, it is easy to understand how a senior manager might read this PowerPoint slide and not realize that it addresses a life-threatening situation.

At many points during its investigation, the Board was surprised to receive similar presentation slides from NASA officials in place of technical reports. The Board views the endemic use of PowerPoint briefing slides instead of technical papers as an illustration of the problematic methods of technical communication at NASA.

**Review Of Test Data Indicates Conservatism for Tile Penetration**

- The existing SOFI on tile test data used to create Crater was reviewed along with STS-107 Southwest Research data
  - Crater overpredicted penetration of tile coating **significantly**
    - Initial penetration to described by normal velocity
      - Varies with volume/mass of projectile (e.g., 200ft/sec for 3cu. In)
    - **Significant** energy is required for the softer SOFI particles to penetrate the relatively hard tile coating
      - Test results do show that it is possible at sufficient mass and velocity
    - Conversely, once tile is penetrated SOFI can cause **significant** damage
      - Minor variations in total energy (above penetration level) can cause **significant** tile damage
  - Flight condition is **significantly** outside of test database
    - Volume of ramp is 1920cu in vs 3 cu in for test

BOEING 2/21/03 6

The vaguely quantitative words "significant" and "significantly" are used 5 times on this slide, with *de facto* meanings ranging from "detectable in largely irrelevant calibration case study" to "an amount of damage so that everyone dies" to "a difference of 640-fold." None of these 5 usages appears to refer to the technical meaning of "statistical significance."

The low resolution of PowerPoint slides promotes the use of compressed phrases like "Tile Penetration." As is the case here, such phrases may well be ambiguous. (The low resolution and large font generate 3 typographic orphans, lonely words dangling on a separate line.)

This vague pronoun reference "it" alludes to *damage to the protective tiles*, which caused the destruction of the Columbia. The slide weakens important material with ambiguous language (sentence fragments, passive voice, multiple meanings of "significant"). The 3 reports were created by engineers for high-level NASA officials who were deciding whether the threat of wing damage required further investigation before the Columbia attempted return. The officials were satisfied that the reports indicated that the Columbia was not in danger, and no attempts to further examine the threat were made. The slides were part of an oral presentation and also were circulated as e-mail attachments.

In this slide the same unit of measure for volume (cubic inches) is shown a different way every time  
**3cu. in**    **1920cu. in**    **3 cu. in**  
 rather than in clear and tidy exponential form **1920 in<sup>3</sup>**. Perhaps the available font cannot show exponents. Shakiness in units of measurement provokes concern. Slides that use hierarchical bullet-outlines here do not handle statistical data and scientific notation gracefully. If PowerPoint is a corporate-mandated format for all engineering reports, then some competent scientific typography (rather than the PP market-pitch style) is essential. In this slide, the typography is so choppy and clunky that it impedes understanding.

The analysis by Dr. Edward Tufte of the slide from the Debris Assessment Team briefing. [SOFI=Spray-On Foam Insulation]



# History As Cause: Columbia and Challenger

The Board began its investigation with two central questions about NASA decisions. Why did NASA continue to fly with known foam debris problems in the years preceding the *Columbia* launch, and why did NASA managers conclude that the foam debris strike 81.9 seconds into *Columbia*'s flight was not a threat to the safety of the mission, despite the concerns of their engineers?

## 8.1 ECHOES OF CHALLENGER

As the investigation progressed, Board member Dr. Sally Ride, who also served on the Rogers Commission, observed that there were “echoes” of *Challenger* in *Columbia*. Ironically, the Rogers Commission investigation into *Challenger* started with two remarkably similar central questions: Why did NASA continue to fly with known O-ring erosion problems in the years before the *Challenger* launch, and why, on the eve of the *Challenger* launch, did NASA managers decide that launching the mission in such cold temperatures was an acceptable risk, despite the concerns of their engineers?

The echoes did not stop there. The foam debris hit was not the single cause of the *Columbia* accident, just as the failure of the joint seal that permitted O-ring erosion was not the single cause of *Challenger*. Both *Columbia* and *Challenger* were lost also because of the failure of NASA's organizational system. Part Two of this report cites failures of the three parts of NASA's organizational system. This chapter shows how previous political, budgetary, and policy decisions by leaders at the White House, Congress, and NASA (Chapter 5) impacted the Space Shuttle Program's structure, culture, and safety system (Chapter 7), and how these in turn resulted in flawed decision-making (Chapter 6) for both accidents. The explanation is about system effects: how actions taken in one layer of NASA's organizational system impact other layers. History is not just a backdrop or a scene-setter. History is cause. History set the *Columbia* and *Challenger* accidents in motion. Although Part Two is separated into chapters and sections to make clear what happened in the political environment, the organization, and managers' and

engineers' decision-making, the three worked together. Each is a critical link in the causal chain.

This chapter shows that both accidents were “failures of foresight” in which history played a prominent role.<sup>1</sup> First, the history of engineering decisions on foam and O-ring incidents had identical trajectories that “normalized” these anomalies, so that flying with these flaws became routine and acceptable. Second, NASA history had an effect. In response to White House and Congressional mandates, NASA leaders took actions that created systemic organizational flaws at the time of *Challenger* that were also present for *Columbia*. The final section compares the two critical decision sequences immediately before the loss of both Orbiters – the pre-launch teleconference for *Challenger* and the post-launch foam strike discussions for *Columbia*. It shows history again at work: how past definitions of risk combined with systemic problems in the NASA organization caused both accidents.

Connecting the parts of NASA's organizational system and drawing the parallels with *Challenger* demonstrate three things. First, despite all the post-*Challenger* changes at NASA and the agency's notable achievements since, the causes of the institutional failure responsible for *Challenger* have not been fixed. Second, the Board strongly believes that if these persistent, systemic flaws are not resolved, the scene is set for another accident. Therefore, the recommendations for change are not only for fixing the Shuttle's technical system, but also for fixing each part of the organizational system that produced *Columbia*'s failure. Third, the Board's focus on the context in which decision making occurred does not mean that individuals are not responsible and accountable. To the contrary, individuals always must assume responsibility for their actions. What it does mean is that NASA's problems cannot be solved simply by retirements, resignations, or transferring personnel.<sup>2</sup>

The constraints under which the agency has operated throughout the Shuttle Program have contributed to both

Shuttle accidents. Although NASA leaders have played an important role, these constraints were not entirely of NASA's own making. The White House and Congress must recognize the role of their decisions in this accident and take responsibility for safety in the future.

## 8.2 FAILURES OF FORESIGHT: TWO DECISION HISTORIES AND THE NORMALIZATION OF DEVIANCE

Foam loss may have occurred on all missions, and left bipod ramp foam loss occurred on 10 percent of the flights for which visible evidence exists. The Board had a hard time understanding how, after the bitter lessons of *Challenger*, NASA could have failed to identify a similar trend. Rather than view the foam decision only in hindsight, the Board tried to see the foam incidents as NASA engineers and managers saw them as they made their decisions. This section gives an insider perspective: how NASA defined risk and how those definitions changed over time for both foam debris hits and O-ring erosion. In both cases, engineers and managers conducting risk assessments continually normalized the technical deviations they found.<sup>3</sup> In all official engineering analyses and launch recommendations prior to the accidents, evidence that the design was not performing as expected was reinterpreted as acceptable and non-deviant, which diminished perceptions of risk throughout the agency.

The initial Shuttle design predicted neither foam debris problems nor poor sealing action of the Solid Rocket Booster joints. To experience either on a mission was a violation of design specifications. These anomalies were signals of potential danger, not something to be tolerated, but in both cases after the first incident the engineering analysis concluded that the design could tolerate the damage. These engineers decided to implement a temporary fix and/or accept the risk, and fly. For both O-rings and foam, that first decision was a turning point. It established a precedent for accepting, rather than eliminating, these technical deviations. As a result of this new classification, subsequent incidents of O-ring erosion or foam debris strikes were not defined as signals of danger, but as evidence that the design was now acting as predicted. Engineers and managers incorporated worsening anomalies into the engineering experience base, which functioned as an elastic waistband, expanding to hold larger deviations from the original design. Anomalies that did not lead to catastrophic failure were treated as a source of valid engineering data that justified further flights. These anomalies were translated into a safety margin that was extremely influential, allowing engineers and managers to add incrementally to the amount and seriousness of damage that was acceptable. Both O-ring erosion and foam debris events were repeatedly "addressed" in NASA's Flight Readiness Reviews but never fully resolved. In both cases, the engineering analysis was incomplete and inadequate. Engineers understood what was happening, but they never understood why. NASA continued to implement a series of small corrective actions, living with the problems until it was too late.<sup>4</sup>

NASA documents show how official classifications of risk were downgraded over time.<sup>5</sup> Program managers designated both the foam problems and O-ring erosion as "acceptable

risks" in Flight Readiness Reviews. NASA managers also assigned each bipod foam event In-Flight Anomaly status, and then removed the designation as corrective actions were implemented. But when major bipod foam-shedding occurred on STS-112 in October 2002, Program management did not assign an In-Flight Anomaly. Instead, it downgraded the problem to the lower status of an "action" item. Before *Challenger*, the problematic Solid Rocket Booster joint had been elevated to a Criticality 1 item on NASA's Critical Items List, which ranked Shuttle components by failure consequences and noted why each was an acceptable risk. The joint was later demoted to a Criticality 1-R (redundant), and then in the month before *Challenger's* launch was "closed out" of the problem-reporting system. Prior to both accidents, this demotion from high-risk item to low-risk item was very similar, but with some important differences. Damaging the Orbiter's Thermal Protection System, especially its fragile tiles, was normalized even before Shuttle launches began: it was expected due to forces at launch, orbit, and re-entry.<sup>6</sup> So normal was replacement of Thermal Protection System materials that NASA managers budgeted for tile cost and turnaround maintenance time from the start.

It was a small and logical next step for the discovery of foam debris damage to the tiles to be viewed by NASA as part of an already existing maintenance problem, an assessment based on experience, not on a thorough hazard analysis. Foam debris anomalies came to be categorized by the reassuring term "in-family," a formal classification indicating that new occurrences of an anomaly were within the engineering experience base. "In-family" was a strange term indeed for a violation of system requirements. Although "in-family" was a designation introduced post-*Challenger* to separate problems by seriousness so that "out-of-family" problems got more attention, by definition the problems that were shifted into the lesser "in-family" category got less attention. The Board's investigation uncovered no paper trail showing escalating concern about the foam problem like the one that Solid Rocket Booster engineers left prior to *Challenger*.<sup>7</sup> So ingrained was the agency's belief that foam debris was not a threat to flight safety that in press briefings after the *Columbia* accident, the Space Shuttle Program Manager still discounted the foam as a probable cause, saying that Shuttle managers were "comfortable" with their previous risk assessments.

From the beginning, NASA's belief about both these problems was affected by the fact that engineers were evaluating them in a work environment where technical problems were normal. Although management treated the Shuttle as operational, it was in reality an experimental vehicle. Many anomalies were expected on each mission. Against this backdrop, an anomaly was not in itself a warning sign of impending catastrophe. Another contributing factor was that both foam debris strikes and O-ring erosion events were examined separately, one at a time. Individual incidents were not read by engineers as strong signals of danger. What NASA engineers and managers saw were pieces of ill-structured problems.<sup>8</sup> An incident of O-ring erosion or foam bipod debris would be followed by several launches where the machine behaved properly, so that signals of danger

safety personnel who remained were ineffective. In the case of *Columbia*, the Board found the same problems were reproduced and for an identical reason: when pressed for cost reduction, NASA attacked its own safety system. The faulty assumption that supported this strategy prior to *Columbia* was that a reduction in safety staff would not result in a reduction of safety, because contractors would assume greater safety responsibility. The effectiveness of those remaining staff safety engineers was blocked by their dependence on the very Program they were charged to supervise. Also, the Board found many safety units with unclear roles and responsibilities that left crucial gaps. Post-*Challenger* NASA still had no systematic procedure for identifying and monitoring trends. The Board was surprised at how long it took NASA to put together trend data in response to Board requests for information. Problem reporting and tracking systems were still overloaded or underused, which undermined their very purpose. Multiple job titles disguised the true extent of safety personnel shortages. The Board found cases in which the same person was occupying more than one safety position – and in one instance at least three positions – which compromised any possibility of safety organization independence because the jobs were established with built-in conflicts of interest.

#### 8.4 ORGANIZATION, CULTURE, AND UNINTENDED CONSEQUENCES

A number of changes to the Space Shuttle Program structure made in response to policy decisions had the unintended effect of perpetuating dangerous aspects of pre-*Challenger* culture and continued the pattern of normalizing things that were not supposed to happen. At the same time that NASA leaders were emphasizing the importance of safety, their personnel cutbacks sent other signals. Streamlining and downsizing, which scarcely go unnoticed by employees, convey a message that efficiency is an important goal. The Shuttle/Space Station partnership affected both programs. Working evenings and weekends just to meet the International Space Station Node 2 deadline sent a signal to employees that schedule is important. When paired with the “faster, better, cheaper” NASA motto of the 1990s and cuts that dramatically decreased safety personnel, efficiency becomes a strong signal and safety a weak one. This kind of doublespeak by top administrators affects people’s decisions and actions without them even realizing it.<sup>26</sup>

Changes in Space Shuttle Program structure contributed to the accident in a second important way. Despite the constraints that the agency was under, prior to both accidents NASA appeared to be immersed in a culture of invincibility, in stark contradiction to post-accident reality. The Rogers Commission found a NASA blinded by its “Can-Do” attitude,<sup>27</sup> a cultural artifact of the Apollo era that was inappropriate in a Space Shuttle Program so strapped by schedule pressures and shortages that spare parts had to be cannibalized from one vehicle to launch another.<sup>28</sup> This can-do attitude bolstered administrators’ belief in an achievable launch rate, the belief that they had an operational system, and an unwillingness to listen to outside experts. The Aerospace Safety and Advisory Panel in a 1985 report told NASA that the vehicle was not operational and NASA should stop

treating it as if it were.<sup>29</sup> The Board found that even after the loss of *Challenger*, NASA was guilty of treating an experimental vehicle as if it were operational and of not listening to outside experts. In a repeat of the pre-*Challenger* warning, the 1999 Shuttle Independent Assessment Team report reiterated that “the Shuttle was not an ‘operational’ vehicle in the usual meaning of the term.”<sup>30</sup> Engineers and program planners were also affected by “Can-Do,” which, when taken too far, can create a reluctance to say that something cannot be done.

How could the lessons of *Challenger* have been forgotten so quickly? Again, history was a factor. First, if success is measured by launches and landings,<sup>31</sup> the machine appeared to be working successfully prior to both accidents. *Challenger* was the 25th launch. Seventeen years and 87 missions passed without major incident. Second, previous policy decisions again had an impact. NASA’s Apollo-era research and development culture and its prized deference to the technical expertise of its working engineers was overridden in the Space Shuttle era by “bureaucratic accountability” – an allegiance to hierarchy, procedure, and following the chain of command.<sup>32</sup> Prior to *Challenger*, the can-do culture was a result not just of years of apparently successful launches, but of the cultural belief that the Shuttle Program’s many structures, rigorous procedures, and detailed system of rules were responsible for those successes.<sup>33</sup> The Board noted that the pre-*Challenger* layers of processes, boards, and panels that had produced a false sense of confidence in the system and its level of safety returned in full force prior to *Columbia*. NASA made many changes to the Space Shuttle Program structure after *Challenger*. The fact that many changes had been made supported a belief in the safety of the system, the invincibility of organizational and technical systems, and ultimately, a sense that the foam problem was understood.

#### 8.5 HISTORY AS CAUSE: TWO ACCIDENTS

Risk, uncertainty, and history came together when unprecedented circumstances arose prior to both accidents. For *Challenger*, the weather prediction for launch time the next day was for cold temperatures that were out of the engineering experience base. For *Columbia*, a large foam hit – also outside the experience base – was discovered after launch. For the first case, all the discussion was pre-launch; for the second, it was post-launch. This initial difference determined the shape these two decision sequences took, the number of people who had information about the problem, and the locations of the involved parties.

For *Challenger*, engineers at Morton-Thiokol,<sup>34</sup> the Solid Rocket Motor contractor in Utah, were concerned about the effect of the unprecedented cold temperatures on the rubber O-rings.<sup>35</sup> Because launch was scheduled for the next morning, the new condition required a reassessment of the engineering analysis presented at the Flight Readiness Review two weeks prior. A teleconference began at 8:45 p.m. Eastern Standard Time (EST) that included 34 people in three locations: Morton-Thiokol in Utah, Marshall, and Kennedy. Thiokol engineers were recommending a launch delay. A reconsideration of a Flight Readiness Review risk



# Part Three

## A Look Ahead

*When it's dark, the stars come out ... The same is true with people. When the tragedies of life turn a bright day into a frightening night, God's stars come out and these stars are families who say although we grieve deeply as do the families of Apollo 1 and Challenger before us, the bold exploration of space must go on. These stars are the leaders in Government and in NASA who will not let the vision die. These stars are the next generation of astronauts, who like the prophets of old said, "Here am I, send me."*

– Brig. Gen. Charles Baldwin, STS-107 Memorial Ceremony at the National Cathedral, February 6, 2003

As this report ends, the Board wants to recognize the outstanding people in NASA. We have been impressed with their diligence, commitment, and professionalism as the agency has been working tirelessly to help the Board complete this report. While mistakes did lead to the accident, and we found that organizational and cultural constraints have worked against safety margins, the NASA family should nonetheless continue to take great pride in their legacy and ongoing accomplishments. As we look ahead, the Board sincerely hopes this report will aid NASA in safely getting back to human space flight.

In Part Three the Board presents its views and recommendations for the steps needed to achieve that goal, of continuing our exploration of space, in a manner with improved safety.

Chapter 9 discusses the near-term, mid-term and long-term implications for the future of human space flight. For the near term, NASA should submit to the Return-to-Flight Task Force a plan for implementing the return-to-flight recommendations. For the mid-term, the agency should focus on: the remaining Part One recommendations, the Part Two recommendations for organizational and cultural changes, and the Part Three recommendation for recertifying the Shuttle for use to 2020 or beyond. In setting the stage for a debate

on the long-term future of human space flight, the Board addresses the need for a national vision to direct the design of a new Space Transportation System.

Chapter 10 contains additional recommendations and the significant "look ahead" observations the Board made in the course of this investigation that were not directly related to the accident, but could be viewed as "weak signals" of future problems. The observations may be indications of serious future problems and must be addressed by NASA.

Chapter 11 contains the recommendations made in Parts One, Two and Three, all issued with the resolve to continue human space flight.





# Recommendations

It is the Board's opinion that good leadership can direct a culture to adapt to new realities. NASA's culture must change, and the Board intends the following recommendations to be steps toward effecting this change.

Recommendations have been put forth in many of the chapters. In this chapter, the recommendations are grouped by subject area with the Return-to-Flight [RTF] tasks listed first within the subject area. Each Recommendation retains its number so the reader can refer to the related section for additional details. These recommendations are not listed in priority order.

## PART ONE – THE ACCIDENT

### Thermal Protection System

- R3.2-1 Initiate an aggressive program to eliminate all External Tank Thermal Protection System debris-shedding at the source with particular emphasis on the region where the bipod struts attach to the External Tank. [RTF]
- R3.3-2 Initiate a program designed to increase the Orbiter's ability to sustain minor debris damage by measures such as improved impact-resistant Reinforced Carbon-Carbon and acreage tiles. This program should determine the actual impact resistance of current materials and the effect of likely debris strikes. [RTF]
- R3.3-1 Develop and implement a comprehensive inspection plan to determine the structural integrity of all Reinforced Carbon-Carbon system components. This inspection plan should take advantage of advanced non-destructive inspection technology. [RTF]
- R6.4-1 For missions to the International Space Station, develop a practicable capability to inspect and effect emergency repairs to the widest possible range of damage to the Thermal Protection System, including both tile and Reinforced Carbon-Carbon, taking advantage of the additional capabilities available when near to or docked at the International Space Station.
- For non-Station missions, develop a comprehensive autonomous (independent of Station) inspection and repair capability to cover the widest possible range of damage scenarios.
- Accomplish an on-orbit Thermal Protection System inspection, using appropriate assets and capabilities, early in all missions.
- The ultimate objective should be a fully autonomous capability for all missions to address the possibility that an International Space Station mission fails to achieve the correct orbit, fails to dock successfully, or is damaged during or after undocking. [RTF]
- R3.3-3 To the extent possible, increase the Orbiter's ability to successfully re-enter Earth's atmosphere with minor leading edge structural sub-system damage.
- R3.3-4 In order to understand the true material characteristics of Reinforced Carbon-Carbon components, develop a comprehensive database of flown Reinforced Carbon-Carbon material characteristics by destructive testing and evaluation.
- R3.3-5 Improve the maintenance of launch pad structures to minimize the leaching of zinc primer onto Reinforced Carbon-Carbon components.
- R3.8-1 Obtain sufficient spare Reinforced Carbon-Carbon panel assemblies and associated support components to ensure that decisions on Reinforced Carbon-Carbon maintenance are made on the basis of component specifications, free of external pressures relating to schedules, costs, or other considerations.

- R3.8-2 Develop, validate, and maintain physics-based computer models to evaluate Thermal Protection System damage from debris impacts. These tools should provide realistic and timely estimates of any impact damage from possible debris from any source that may ultimately impact the Orbiter. Establish impact damage thresholds that trigger responsive corrective action, such as on-orbit inspection and repair, when indicated.

### Imaging

- R3.4-1 Upgrade the imaging system to be capable of providing a minimum of three useful views of the Space Shuttle from liftoff to at least Solid Rocket Booster separation, along any expected ascent azimuth. The operational status of these assets should be included in the Launch Commit Criteria for future launches. Consider using ships or aircraft to provide additional views of the Shuttle during ascent. [RTF]
- R3.4-2 Provide a capability to obtain and downlink high-resolution images of the External Tank after it separates. [RTF]
- R3.4-3 Provide a capability to obtain and downlink high-resolution images of the underside of the Orbiter wing leading edge and forward section of both wings' Thermal Protection System. [RTF]
- R6.3-2 Modify the Memorandum of Agreement with the National Imagery and Mapping Agency to make the imaging of each Shuttle flight while on orbit a standard requirement. [RTF]

### Orbiter Sensor Data

- R3.6-1 The Modular Auxiliary Data System instrumentation and sensor suite on each Orbiter should be maintained and updated to include current sensor and data acquisition technologies.
- R3.6-2 The Modular Auxiliary Data System should be redesigned to include engineering performance and vehicle health information, and have the ability to be reconfigured during flight in order to allow certain data to be recorded, telemetered, or both as needs change.

### Wiring

- R4.2-2 As part of the Shuttle Service Life Extension Program and potential 40-year service life, develop a state-of-the-art means to inspect all Orbiter wiring, including that which is inaccessible.

### Bolt Catchers

- R4.2-1 Test and qualify the flight hardware bolt catchers. [RTF]

### Closeouts

- R4.2-3 Require that at least two employees attend all final closeouts and intertank area hand-spraying procedures. [RTF]

### Micrometeoroid and Orbital Debris

- R4.2-4 Require the Space Shuttle to be operated with the same degree of safety for micrometeoroid and orbital debris as the degree of safety calculated for the International Space Station. Change the micrometeoroid and orbital debris safety criteria from guidelines to requirements.

### Foreign Object Debris

- R4.2-5 Kennedy Space Center Quality Assurance and United Space Alliance must return to the straightforward, industry-standard definition of "Foreign Object Debris" and eliminate any alternate or statistically deceptive definitions like "processing debris." [RTF]

## PART TWO – WHY THE ACCIDENT OCCURRED

### Scheduling

- R6.2-1 Adopt and maintain a Shuttle flight schedule that is consistent with available resources. Although schedule deadlines are an important management tool, those deadlines must be regularly evaluated to ensure that any additional risk incurred to meet the schedule is recognized, understood, and acceptable. [RTF]

### Training

- R6.3-1 Implement an expanded training program in which the Mission Management Team faces potential crew and vehicle safety contingencies beyond launch and ascent. These contingencies should involve potential loss of Shuttle or crew, contain numerous uncertainties and unknowns, and require the Mission Management Team to assemble and interact with support organizations across NASA/Contractor lines and in various locations. [RTF]

## Organization

- R7.5-1 Establish an independent Technical Engineering Authority that is responsible for technical requirements and all waivers to them, and will build a disciplined, systematic approach to identifying, analyzing, and controlling hazards throughout the life cycle of the Shuttle System. The independent technical authority does the following as a minimum:
- Develop and maintain technical standards for all Space Shuttle Program projects and elements
  - Be the sole waiver-granting authority for all technical standards
  - Conduct trend and risk analysis at the sub-system, system, and enterprise levels
  - Own the failure mode, effects analysis and hazard reporting systems
  - Conduct integrated hazard analysis
  - Decide what is and is not an anomalous event
  - Independently verify launch readiness
  - Approve the provisions of the recertification program called for in Recommendation R9.1-1.

The Technical Engineering Authority should be funded directly from NASA Headquarters, and should have no connection to or responsibility for schedule or program cost.

- R7.5-2 NASA Headquarters Office of Safety and Mission Assurance should have direct line authority over the entire Space Shuttle Program safety organization and should be independently resourced.
- R7.5-3 Reorganize the Space Shuttle Integration Office to make it capable of integrating all elements of the Space Shuttle Program, including the Orbiter.

## PART THREE – A LOOK AHEAD

### Organization

- R9.1-1 Prepare a detailed plan for defining, establishing, transitioning, and implementing an independent Technical Engineering Authority, independent safety program, and a reorganized Space Shuttle Integration Office as described in R7.5-1, R7.5-2, and R7.5-3. In addition, NASA should submit annual reports to Congress, as part of the budget review process, on its implementation activities. [RTF]

## Recertification

- R9.2-1 Prior to operating the Shuttle beyond 2010, develop and conduct a vehicle recertification at the material, component, subsystem, and system levels. Recertification requirements should be included in the Service Life Extension Program.

### Closeout Photos/Drawing System

- R10.3-1 Develop an interim program of closeout photographs for all critical sub-systems that differ from engineering drawings. Digitize the closeout photograph system so that images are immediately available for on-orbit troubleshooting. [RTF]
- R10.3-2 Provide adequate resources for a long-term program to upgrade the Shuttle engineering drawing system including:
- Reviewing drawings for accuracy
  - Converting all drawings to a computer-aided drafting system
  - Incorporating engineering changes

R4.2-1 Test and qualify the flight hardware bolt catchers.

The fault tree review **brought to light a significant problem with the Solid Rocket Booster bolt catchers.** (p. 86)

Two "bolt catchers" on the External Tank each trap the upper half of a fired separation bolt, while the lower half stays attached to the Solid Rocket Booster. As a result, both halves are kept from flying free of the assembly and potentially hitting the Orbiter. Bolt catchers have a domed aluminum cover containing an aluminum honeycomb matrix that absorbs the fired bolt's energy. The two upper bolt halves and their respective catchers subsequently remain connected to the External Tank, which burns up on re-entry, while the lower halves stay with the Solid Rocket Boosters that are recovered from the ocean. (p. 86)

**... the configuration of the bolt catchers used on Shuttle missions differs in important ways from the design used in initial qualification tests.** First, the attachments that currently hold bolt catchers in place use bolts threaded into inserts rather than through-bolts. Second, the test design included neither the Super Lightweight Ablative material applied to the bolt catcher apparatus for thermal protection, nor the aluminum honeycomb configuration currently used. Also, during these initial tests, temperature and pressure readings for the bolt firings were not recorded. (pp. 86-87)

**The flight configuration was validated using extrapolated test data and redesign specifications rather than direct testing.** This means that NASA's rationale for considering bolt catchers to be safe for flight is based on limited data from testing 24 years ago on a model that differs significantly from the current design. (p. 87)

**Due to these testing deficiencies, the Board recognized that bolt catchers could have played a role in damaging Columbia's left wing.** (p. 87)

Although bolt catchers can be neither definitively excluded nor included as a potential cause of left wing damage to Columbia, the impact of such a large object would likely have registered on the Shuttle stack's sensors. The indefinite data at the time of Solid Rocket Booster separation, in tandem with overwhelming evidence related to the foam debris strike, leads **the Board to conclude that bolt catchers are unlikely to have been involved in the accident.** (p. 88)