AD ASTRA PER ASPERA —

The audacious rescue plan that might have saved space shuttle *Columbia*

The untold story of the rescue mission that could have been NASA's finest hour.

LEE HUTCHINSON - 2/1/2016, 7:30 AM

Lee Hutchinson / NASA / NOAA

Enlarge / What might have been.

February 1, 2023: One of the most tragic events in the history of space exploration is the loss of the space shuttle Columbia and all seven of her crew on February 1, 2003—a tragedy made worse because it didn't have to happen. But just as it is human nature to look to the future and wonder what might be, so too is it in our nature to look at the past and wonder, "what if?" Today, on the twentieth anniversary of the event, Ars is re-publishing our detailed 2014 examination of the biggest Columbia "what if" of all: what if NASA had recognized the damage to the orbiter while the mission was still in progress? Could anything have been done to save the crew?

> ***If we die, we want people to accept it. We are in a risky business, and we hope that if anything happens to us, it will not delay the program. The conquest of space is worth the risk of life.***
>
> —Astronaut Gus Grissom, 1965

> ***It is important to note at the outset that Columbia broke up during a phase of flight that, given the current design of the Orbiter, offered no possibility of crew survival.***
>
> —Columbia Accident Investigation Board Report

At 10:39 Eastern Standard Time on January 16, 2003, space shuttle *Columbia* lifted off from pad 39A at the Kennedy Space Center in Florida. Just under a minute and a half later, at 81.7 seconds after launch, a chunk of insulating foam tore free from the orange external tank and smashed into the leading edge of the orbiter's left wing at a relative velocity of at least 400 miles per hour (640 km/h). *Columbia* continued to climb toward orbit.

The foam strike was not observed live. Only after the shuttle was orbiting Earth did NASA's launch imagery review reveal that the wing had been hit. Foam strikes during launch were not uncommon events, and shuttle program managers elected not to take on-orbit images of *Columbia* to visually assess any potential damage. Instead, NASA's Debris Assessment Team mathematically modeled the foam strike but could not reach any definitive conclusions about the state of the shuttle's wing. The mission continued.

In reality, the impact shattered at least one of the crucial reinforced carbon-carbon heat shield panels that lined the edge of the wing, leaving a large hole in the brittle ceramic material. Sixteen days later, as *Columbia* re-entered the atmosphere, superheated plasma entered the orbiter's structure through the hole in the wing and the shuttle began to disintegrate.

At Mission Control in Houston, the flight controllers monitoring *Columbia*'s descent began to notice erratic telemetry readings coming from the shuttle, and then all voice and data contact with the orbiter was lost. Controllers continued to hope that they were merely looking at instrumentation failures, even as evidence mounted that a catastrophic event had taken place. Finally, at 9:12 Eastern Time, re-entry Flight Director LeRoy Cain keyed his communications loop and called out a rarely heard order: "Lock the doors."

It was an acknowledgement that the worst had happened; the mission was now in "contingency" mode. The control room was sealed off, and each flight controller began carefully preserving his or her console's data.

Columbia was gone, and all seven of her crew had been killed. NASA refers to this most rare and catastrophic of events as an LOCV—"Loss of Crew and Vehicle."

Frozen

> ***Columbia is lost. There are no survivors.***
>
> —President George W. Bush in a national address, 14:04 EST, February 1, 2003

The world of human space flight paused—first to mourn, then to discover what had happened. Congress laid that responsibility on the combined shoulders of the Columbia Accident Investigation Board (referred to, in typical NASA acronym-dependent style, as "the CAIB" or just "CAIB," which rhymes with "Gabe"). In the months after *Columbia*, the CAIB stretched its investigative fingers all through NASA and its supporting contractors.

My own memories of the time immediately following the accident are dominated by images of somber meetings and frantic work. I was a junior system administrator at Boeing in Houston, and because we supported the shuttle program, we had to locate and send cases and cases of backup tapes—containing everything that happened on every server in our data center during the mission—over to NASA for analysis.

In August 2003, the CAIB issued its final report. Behind the direct cause of the foam strike, the report leveled damning critiques at NASA's pre- and post-launch decision-making, painting a picture of an agency dominated by milestone-obsessed middle management. That focus on narrow, group-specific work and reporting, without a complementary focus on cross-department integration and communication, contributed at least as much to the loss of the shuttle as did the foam impact. Those accusations held a faint echo of familiarity—many of them had been raised 17 years earlier by the Rogers Commission investigating *Challenger's* destruction.

In the end, *Columbia*'s loss ended not only lives but also careers at all levels of NASA. A number of prominent shuttle program managers were reassigned. It is likely that *Columbia*'s destruction factored heavily into the resignation of NASA Administrator Sean O'Keefe. Many involved with the mission—including many still working at NASA—to this day struggle with post-traumatic stress and survivor's guilt. All pending shuttle missions were put on hold, and *Columbia*'s three surviving sister ships—*Discovery*, *Atlantis*, and *Endeavour*—were grounded.

NASA looked inward, and we wondered if the orbiters would ever fly again.

A path not taken

> ***To put the decisions made during the flight of STS-107 into perspective, the Board asked NASA to determine if there were options for the safe return of the STS-107 crew.***
>
> —Columbia Accident Investigation Board Report

That's the way events actually unfolded. But imagine an alternate timeline for the *Columbia* mission in which NASA quickly realized just how devastating the foam strike had been. Could the *Columbia* astronauts have been safely retrieved from orbit?

The CAIB had the same question during the writing of its report, so it asked NASA to develop a theoretical repair and rescue plan for *Columbia* "based on the premise that the wing damage events during launch were recognized early during the mission." The result was a remarkable document, which appears at the end of the report as Appendix D.13. It carries the low-key title "STS-107 In-Flight Options Assessment," but the scenario it outlines would have pushed NASA to its absolute limits as the agency mounted the most dramatic space mission of all time.

NASA planners had one fortuitous ace in the hole that made the whole plan possible: while *Columbia*'s STS-107 mission was in progress, *Atlantis* was already undergoing preparation for flight as STS-114, scheduled for launch on March 1. As *Columbia* thundered into orbit, the younger shuttle was staged in Orbital Processing Facility 1 (OPF-1) at the Kennedy Space Center. Her three main engines had already been installed, but she didn't yet have a remote manipulator arm or a payload for her payload bay. Two more weeks of refurbishment and prep work remained before *Atlantis* would be wheeled across the space center to the enormous Vehicle Assembly Building and hoisted up for attachment to an external tank and a pair of solid rocket boosters.

NASA

Enlarge / *Endeavour* undergoes processing at OPF-2. *Atlantis* was in a similar state while *Columbia* was flying its final mission.

So an in-orbit rescue was at least *feasible*—but making a shuttle ready to fly is an incredibly complicated procedure involving millions of discrete steps. In order to pull *Atlantis'* launch forward, mission planners had to determine which steps (if any) could be safely skipped without endangering the rescue crew.

The desperate race

> ***The scenarios were to assume that a decision to repair or rescue the Columbia crew would be made quickly, with no regard to risk.***
>
> —Columbia Accident Investigation Board Report (Appendix D.13)

But even before those decisions could be made, NASA had to make another assessment—how long did it have to mount a rescue? In tallying *Columbia's* supplies, NASA mission planners realized that the most pressing supply issue for the astronauts wasn't running *out* of something like air or water but accumulating *too much* of something: carbon dioxide.

Weight is a precious commodity for spacecraft. Every gram of mass boosted up into orbit must be paid for with fuel, and adding fuel adds weight that must also be paid for in more fuel (this spiral of mass-begets-fuel-begets-mass is often referred to as the tyranny of the rocket equation). Rather than carrying up spare "air," spacecraft launch with a mostly fixed volume of internal air, which they recycle by adding back component gasses. The space shuttle carries supplies of liquid oxygen and liquid nitrogen, which are turned into gas and cycled into the cabin's air to maintain a 78 percent nitrogen/21 percent oxygen mixture, similar to Earth's atmosphere. The crew exhales carbon dioxide, though, and that carbon dioxide must be removed from the air.

To do this, the shuttle's air is filtered through canisters filled with lithium hydroxide (LiOH), which attaches to carbon dioxide molecules to form lithium carbonate crystals (Li_2CO_3), thus sequestering the toxic carbon dioxide. These canisters are limited-use items, each containing a certain quantity of lithium hydroxide; *Columbia* was equipped with 69 of them.

How long those 69 canisters would last proved difficult to estimate, though, because there isn't a lot of hard data on how much carbon dioxide the human body can tolerate in microgravity. Standard mission operation rules dictate that the mission be aborted if CO_2 levels rise above a partial pressure of 15 mmHg (about two percent of the cabin air's volume), and mission planners believed they could stretch *Columbia*'s LiOH canister supply to cover a total of 30 days of mission time without breaking that CO_2 threshold. However, doing so would require the crew to spend 12 hours of each day doing as little as possible—sleeping, resting, and doing everything they could to keep their metabolic rates low.

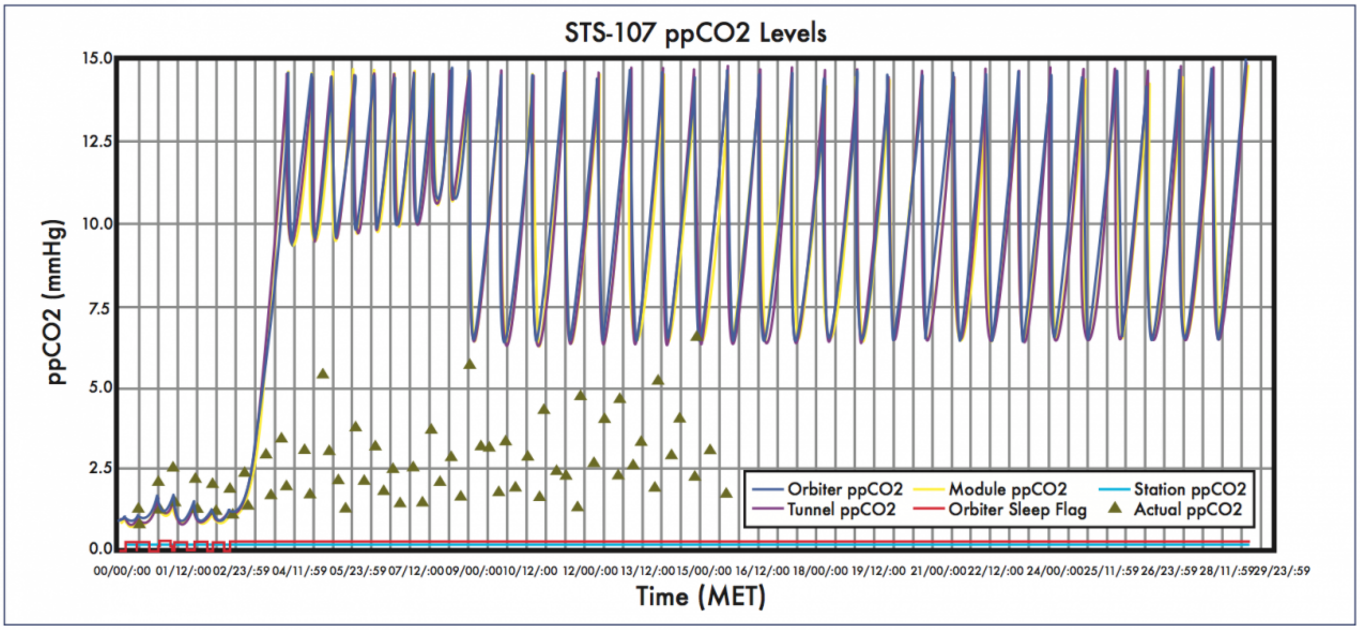

Figure 2. ppCO2 plot with 12 hours of Crew Sleep.

NASA / CAIB Report, Appendix D.13

Enlarge

If the crew couldn't sustain that low rate of activity, NASA flight surgeons believed that allowing the CO_2 content to rise to a partial pressure of 26.6 mmHg (about 3.5 percent cabin air volume) "would not produce any long-term effects on the health of the crewmembers." This would enable the crew to function on a more "normal" 16-hour/8-hour wake/sleep cycle, but at the cost of potential physiological deficits; headaches, fatigue, and other problems related to the high CO_2 levels would have started to manifest very quickly.

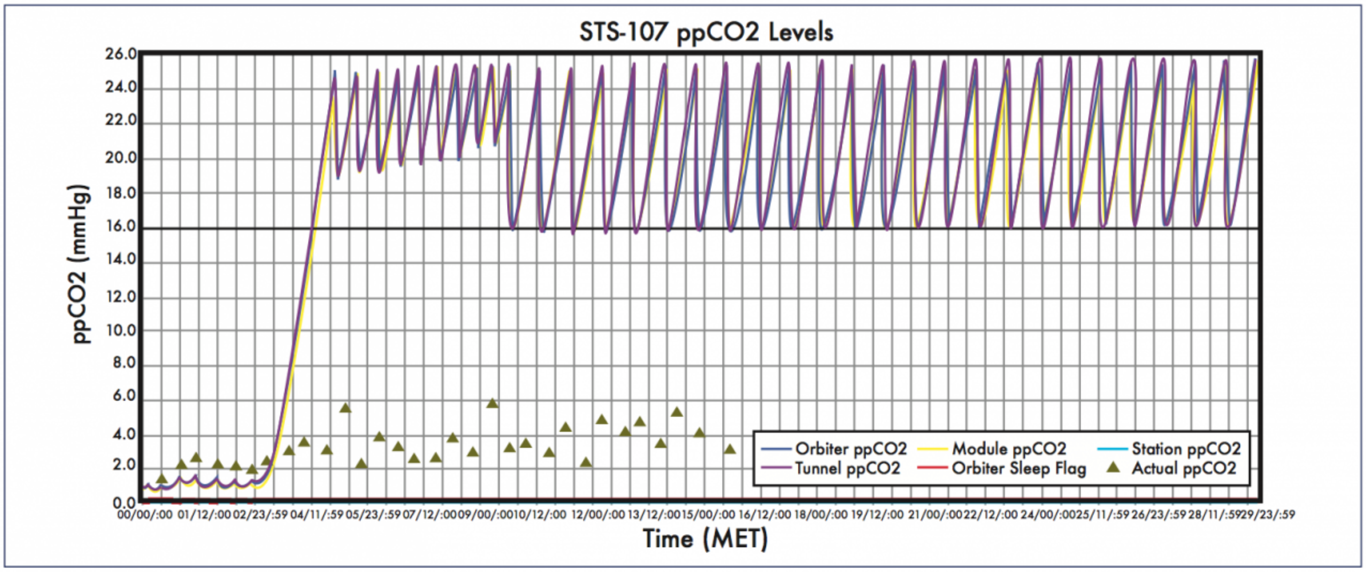

NASA / CAIB Report, Appendix D.13

Figure 1. ppCO2 plot with 8 hours of Crew Sleep.

Enlarge

After the carbon dioxide scrubbers, the next most limited consumable was oxygen. *Columbia*'s liquid oxygen supplies were used not only to replenish breathing gas for the crew but also to generate power in the shuttle's fuel cells (which combined oxygen with hydrogen to produce both energy and potable water). The amount of liquid oxygen on board could be stretched past the CO_2 scrubbers' 30-day mark by drastically cutting down *Columbia*'s power draw.

The remaining three consumable categories consisted of food, water, and propellant. Assuming that the crew would be moving minimally, food and water could stretch well beyond the 30-day limit imposed by the LiOH canisters. To preserve propellant, the orbiter would be placed into an attitude needing minimal fuel to maintain.

Exactly *when* the crew of *Columbia* would enact these power- and oxygen-saving measures depended on a short decision tree. In the scenario we're walking through, the assumption is that NASA determined on Flight Day 2 (January 17) that the foam strike had caused some damage, followed by at least another day to gather images of *Columbia* using "national assets" like ground-based telescopes and other space-based sources (i.e., spy satellites) under the control of USSTRATCOM.

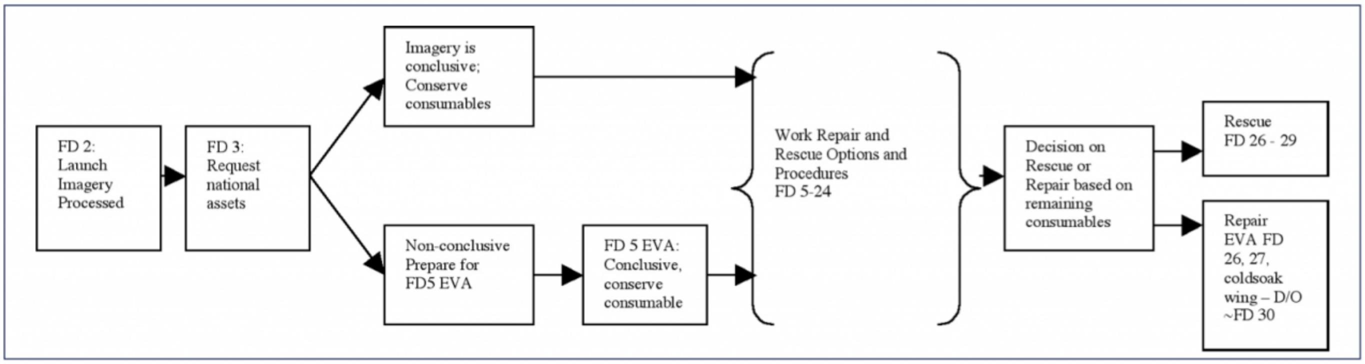

NASA / CAIB Report, Appendix D.13

Figure 3. Decision path timeline.

Enlarge / "FD" is short for "Flight Day."

If that imagery positively identified damage, *Columbia* would immediately enter power-down mode; if the images didn't show anything conclusive, the crew would conduct an EVA (extra-vehicular activity—a spacewalk) to visually assess the damage to the wing, then power things down.

In either case, Flight Day 3 would mark the start of many sleepless nights for many people.

No do-overs, no mistakes

> ***This rescue was considered challenging but feasible.***
>
> —Columbia Accident Investigation Board Report

Planning the inspection EVA would have taken most of Flight Day 4 (January 19), but the hard deadline of the lithium hydroxide canisters remained set at Flight Day 30 (February 15) regardless of what happened on the ground. Work would simultaneously have had to begin at the Kennedy Space Center to accelerate the processing of *Atlantis*.

"Accelerate" is a prosaic word for the herculean effort that would have been needed. Activities that normally take place across weeks or months would have to happen in hours or days. Civil servants and contractors at KSC would have to begin 24/7 shift work, keeping the lights on and the process running every hour of every day, for a minimum of *21 days*, to power *Atlantis* through checkout and make her ready to launch.

Three unceasing, brutal weeks of 24/7 shift work—and that's with absolutely no margin factored in for errors or failures. The Orbital Processing Facility team, the Vehicle Assembly Building team, and the Launch Complex 39 pad team would have had to get every one of the millions of steps right, and every component of *Atlantis* would have had to function perfectly the very first time, or it would all be wasted.

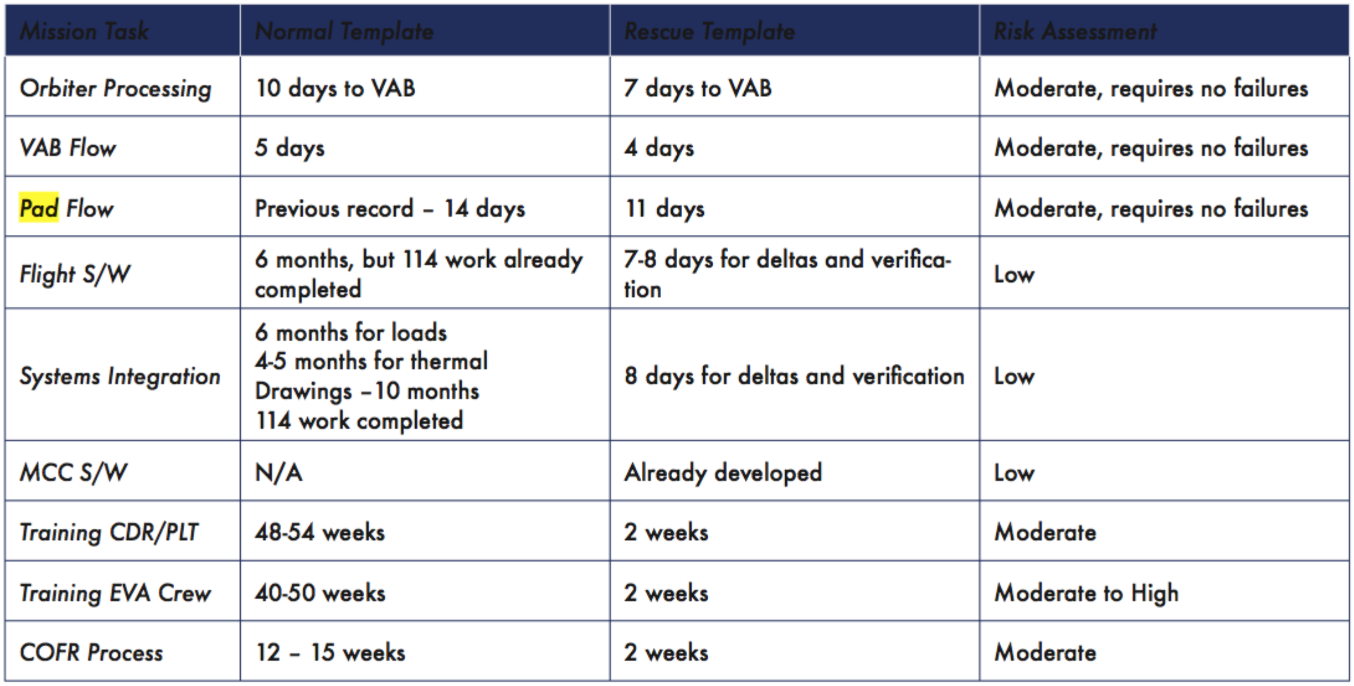

Mission Task	Normal Template	Rescue Template	Risk Assessment
Orbiter Processing	10 days to VAB	7 days to VAB	Moderate, requires no failures
VAB Flow	5 days	4 days	Moderate, requires no failures
Pad Flow	Previous record – 14 days	11 days	Moderate, requires no failures
Flight S/W	6 months, but 114 work already completed	7-8 days for deltas and verification	Low
Systems Integration	6 months for loads 4-5 months for thermal Drawings –10 months 114 work completed	8 days for deltas and verification	Low
MCC S/W	N/A	Already developed	Low
Training CDR/PLT	48-54 weeks	2 weeks	Moderate
Training EVA Crew	40-50 weeks	2 weeks	Moderate to High
COFR Process	12 – 15 weeks	2 weeks	Moderate

NASA / CAIB Report, Appendix D.13

Enlarge / **A rescue mission would require preparing a shuttle for launch far faster than had ever been done before.**

So many things would have to happen. First, *Atlantis'* computers would have to be reprogrammed to accommodate the changes in the mission. Fortunately, the flight software developed for STS-114's International Space Station (ISS) rendezvous could be adapted to instead rendezvous with *Columbia*, though most of the specific rendezvous parameters would have to be altered. The changes would be uploaded to *Atlantis'* computers during the DOLILU—the Day of Launch Input Load Update, the standard last-minute software update that shuttles on the pad receive two hours prior to launch. Usually, DOLILU loads include flight control updates to accommodate the day's observed weather patterns, but this particular DOLILU load would change the entire flight profile. It would be the largest on-pad software update ever attempted.

In order to push *Atlantis* through processing in time, a number of standard checks would have to be abandoned. The expedited OPF processing would get *Atlantis* into the Vehicle Assembly Building in just six days, and the 24/7 prep work would then shave an additional day off the process of getting *Atlantis* mated to external tank and boosters. After only four days in the Vehicle Assembly Building, one of the two Crawler-Transporters would haul *Atlantis* out to Launch Complex 39, where she would stage on either Pad A or Pad B on Flight Day 15—January 30.

NASA / WikiMedia Commons

Enlarge / **Under troubled skies, *Atlantis* makes her way out to the pad atop one of the Crawler-Transporters to embark on STS-129.**

Once on the pad, the final push to launch would begin. There would be no practice countdown for the astronauts chosen to fly the mission, nor would there be extra fuel leak tests. Prior to this launch, the shortest time a shuttle had spent on the launch pad was 14 days; the pad crews closing out *Atlantis* would have only 11 days to get her ready to fly.

Even as workers at Cape Canaveral frantically tried to beat the clock, more work had to happen at the Johnson Space Center in Houston: *Atlantis* still needed a crew.

The right stuff

> ***[I]t would be important to have a high degree of confidence in the astronauts' ability to quickly adapt to the micro-gravity environment.***
>
> —Columbia Accident Investigation Board Report (Appendix D.13)

Columbia carried seven astronauts, who by Flight Day 15 would be halfway through their unexpectedly extended 30-day mission. This presented a problem for NASA: *Atlantis* would need her own crew in order to launch and rendezvous with *Columbia*, but space shuttles were only designed to accommodate five to seven astronauts. When *Atlantis* returned, she would carry not only the astronauts she launched with, but also *Columbia*'s rescued crew—so, to minimize crowding, what was the minimal crew count *Atlantis* could get away with at launch?

After analysis, it was determined that *Atlantis* would need a minimum crew of four. A two-person pilot and commander team would be required to actually fly the rendezvous and actively keep station with *Columbia*—which NASA estimated would mean at least eight or nine hours of manual flying (and potentially much more than that). Another two-person team would be required to don suits and perform the rescue EVA tasks—tasks which NASA would have had to design from scratch.

As with every other piece of the rescue, there was no room for error, and there would be no second chances. *Atlantis* would therefore be launched with an all-veteran crew, with selection for the mission biased heavily toward astronauts who demonstrated fast adaptation to microgravity (there was no time to be space-sick) and high aptitude at EVA or rendezvous. The report names no names, but it does indicate that an assessment revealed a pool of nine EVA candidates, seven command candidates, and seven pilot candidates available in January 2003 whom NASA felt could have undertaken the mission.

FURTHER READING
Swimming with spacemen: training for spacewalks at NASA's giant pool

The four astronauts chosen to fly *Atlantis* would have faced an extraordinarily compressed training schedule—and also a tremendous amount of professional and personal pressure. The tight timeline would mean that the two *Atlantis* astronauts selected to actually spacewalk between the shuttles for the rescue EVA would likely be training underwater at NASA's Neutral Buoyancy Lab almost every single day of the two weeks, breaking the entire multi-hour spacewalk up into tiny component maneuvers and procedures and walking through each to commit them to memory. Simultaneously, the two astronauts selected to pilot the shuttle would have spent that time in the large motion-base simulator in Building 9 at the Johnson Space Center, working through every moment of the rendezvous, station-keeping, and landing from start to finish.

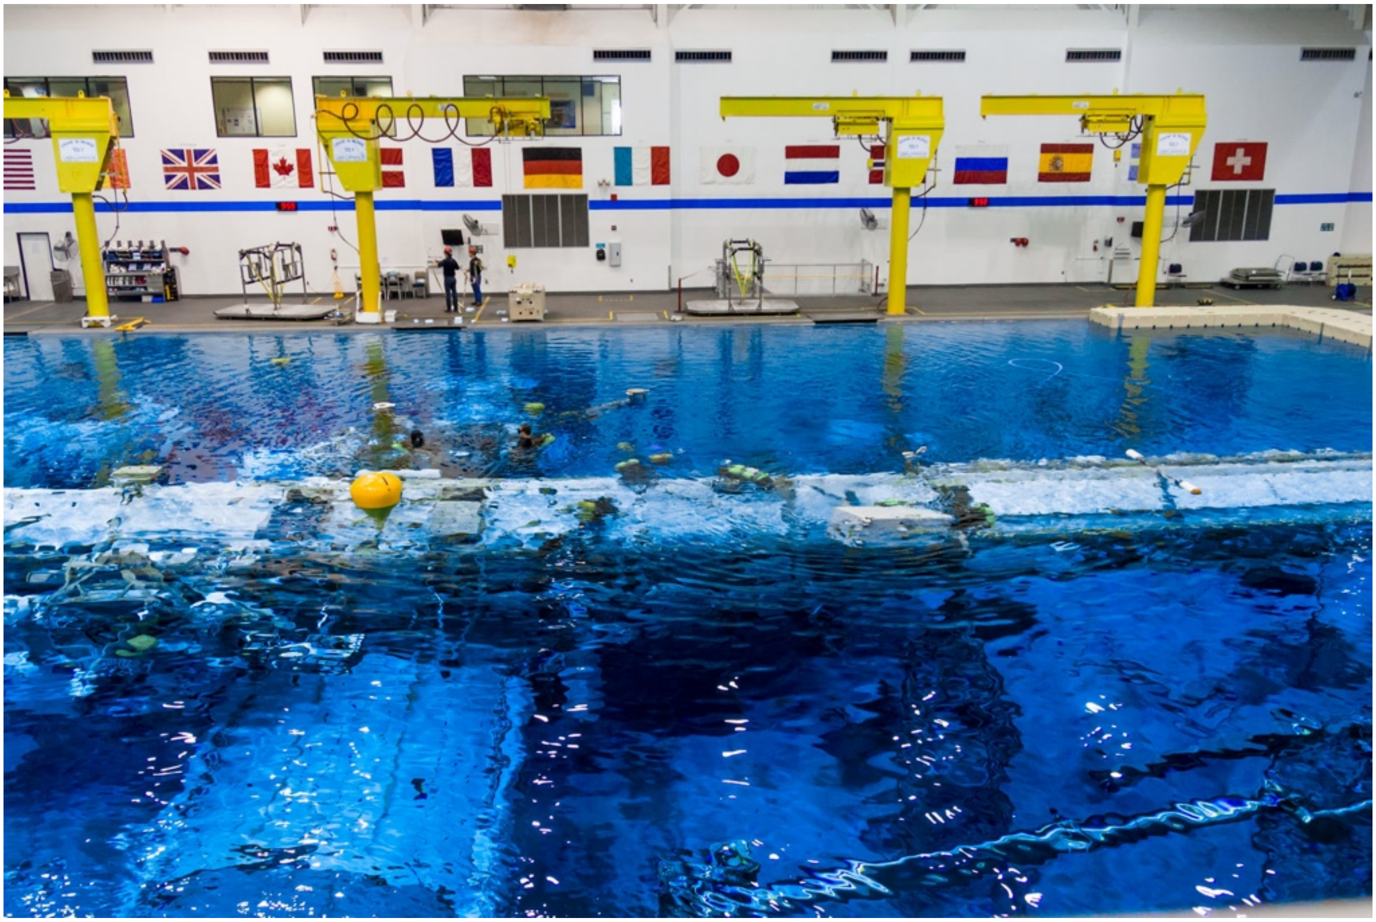
Steven Michael

Enlarge / **Looking down into NASA's Neutral Buoyancy Lab pool from one of the test director control rooms.**

It's also certain that the media would have exerted its own tremendous pressure, attempting to thrust cameras and lights into every corner of the preparation—as much as they would be allowed to do so, anyway. "Space disaster" and "rescue mission" are golden ratings words. Clear Lake in Houston and Cape Canaveral in Florida would have been swarmed with TV trucks; the Johnson Space Center sign on historic NASA Rd 1 would likely have been a constant backdrop on TV news both local and national.

And throughout the frantic weeks on the ground, *Columbia*'s crew would wait.

Slow time

> ***This powerdown would have supported only the most basic vehicle control and crew support and communication equipment.***
>
> —Columbia Accident Investigation Board Report (Appendix D.13)

While work on the ground would proceed in a controlled frenzy, time on *Columbia* would lengthen and draw out in slow misery. The crew would have potentially undergone a brief flurry of activity if they needed to do an EVA to confirm the damage to the orbiter's left wing; additionally, they would have needed to maneuver *Columbia* into a tail-first "gravity gradient" attitude so that the Earth's pull on the shuttle's empennage would keep the orbiter's orientation fixed relative to Earth without the need to expend any propellant. After that, though, the stranded crew could do very little other than wait and try not to move or breathe too much.

NASA / WikiMedia Commons

Enlarge / STS-107 mission specialists Laurel B. Clark, Rick D. Husband, and Kalpana Chawla relaxing in their bunks on *Columbia*'s middeck.

The crew wouldn't even be able to watch their own rescue's TV coverage, because the orbiter would be in a tightly restricted low-power mode in order to conserve its energy. Appendix D.13 includes a description of what systems would be shut down, and among them are "all cameras, camera heaters, TV monitors, and video equipment."

An oft-asked question is whether or not *Columbia* could have docked with the ISS, which would have had consumables to spare. There are numerous reasons why this would not have been possible, but the overriding one comes down to simple physics: *Columbia* would have had to execute what is known in orbital mechanics terminology as a "plane change" maneuver—applying thrust perpendicular to her orbital track to shift orbit to match the ISS' more tilted inclination. Plane change maneuvers require tremendous amounts of energy—in some cases, even more energy than was required to launch the spacecraft in the first place. Appendix D.13 dismisses the possibility of an ISS rendezvous with just two sentences:

> *Columbia*'s 39 degree orbital inclination could not have been altered to the ISS 51.6 degree inclination without approximately 12,600 ft/sec of translational capability. *Columbia* had 448 ft/sec of propellant available.

The crew would be playing the long game, carefully conserving resources for the burst of activity that would have to occur at the end of the mission. As previously noted, the primary consumable of consequence would be the carbon dioxide scrubbers, so rest and sleep would have been the crew's main mission. *Columbia*'s orbital period would mean that during their quiet exile, the crew would see more than 300 sunrises break over the curving lip of the world, spread across weeks of endless drifting purgatory.

Russian roulette

> ***This new risk to the Orbiter would weigh heavily in the decision process on launching another shuttle and crew.***
>
> —Columbia Accident Investigation Board Report (Appendix D.13)

Appendix D.13 is written under the assumption that the damage to *Columbia*'s wing was recognized and acted upon, but that is actually the first of *two* major assumptions underlying the rescue mission. The second assumption has its own set of enormous issues: given that *Columbia* was disabled by a foam strike, NASA would have to be willing to subject *Atlantis* to the exact same risk.

The obvious unsettling question here is whether or not there was anything NASA could do in the near term to prevent *Atlantis* from being disabled by the same type of foam impact—and, given the nature of the strike, the answer is "no."

The foam chunk that sheared off of *Columbia*'s external tank was part of what's called the left "bipod ramp," one of two hand-sculpted structures flanking the large bipod struts that secure the orbiter's nose to the forward part of the external tank. To form the bipod ramps, orange BX-250 insulation is sprayed over the fittings that attach the bipod to the external tank. It's allowed to dry, then it's shaved by hand into wedges that cover the fitting elements. Coupled with a layer of ablative materials atop the fittings, the foam ramps both protect the attachment points from heat during launch and also sheath them in an aerodynamic shape.

And, as it turns out, bipod ramps had broken off *six times* before STS-107.

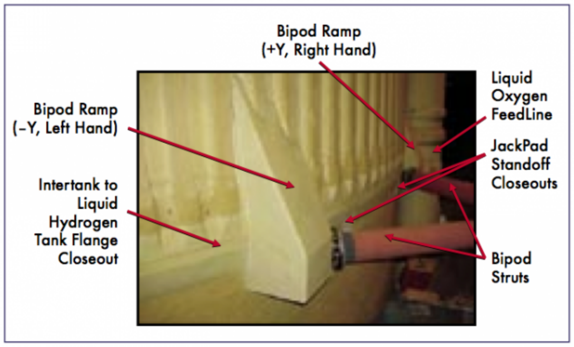

Figure 3.2-2. The exterior of the left bipod attachment area showing the foam ramp that came off during the ascent of STS-107.

NASA / CAIB Report, Appendix D.13

Enlarge

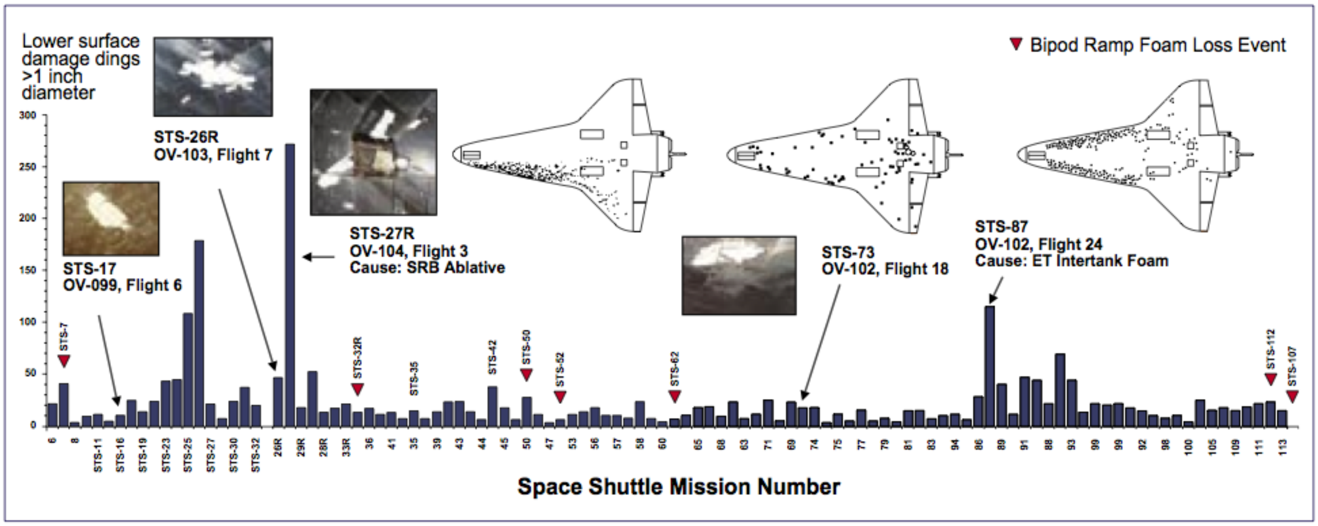

Figure 6.1-6. This chart shows the number of dings greater than one inch in diameter on the lower surface of the Orbiter after each mission from STS-6 through STS-113. Flights where the bipod ramp foam is known to have come off are marked with a red triangle.

NASA / CAIB Report, Appendix D.13

Enlarge

Much of the CAIB report is given over to discussing the specifics of the external tank's insulating foam—what it's made of, how that material performs, and how often foam has sheared off of the ET and impacted with an orbiter. What is clear from the report is that the STS-107 foam strike was not a unique event—it was a relatively common occurrence that in this particular instance happened at precisely the right (or wrong) time to cause catastrophic damage to one of the very few things on the shuttle without any form of redundancy.

Page: 1 2 3 4 5 6 7 8 9 10 Next →

READER COMMENTS 507

LEE HUTCHINSON

Lee is the Senior Technology Editor, and oversees story development for the gadget, culture, IT, and video sections of Ars Technica. A long-time member of the Ars OpenForum with an extensive background in enterprise storage and security, he lives in Houston.

WATCH

Biomarkers, from diagnosis to treatment

Biomarkers, from diagnosis to treatment

To find cancer before it strikes, look for the molecular clues

Biomarkers, from diagnosis to treatment

Scott Manley Reacts To His Top 1000 YouTube Comments

Teach the Controversy: Dowsing

⊕ More videos

← PREVIOUS STORY NEXT STORY →

Related Stories

Today on Ars

The $1,600 Dyson 360 Vis Nav promises to be the world’s most powerful robovac

Woman with untreated TB still on the lam three months after arrest warrant

Study: People think undermining democracy is ok if others do it first

Google bans Downloader app after TV firms complain it can load a pirate website

Adobe Photoshop’s new “Generative Fill” AI tool lets you manipulate photos with text

Meta has no choice but to sell Giphy at $262M loss to Shutterstock

Unexpected 3DS update breaks many common homebrew hacking methods

Built-in ChatGPT-driven Copilot will transform Windows 11 starting in June

STORE
SUBSCRIBE
ABOUT US
RSS FEEDS
VIEW MOBILE SITE

CONTACT US
STAFF
ADVERTISE WITH US
REPRINTS

NEWSLETTER SIGNUP

Join the Ars Orbital Transmission mailing list to get weekly updates delivered to your inbox. Sign me up →