National Aeronautics and Space Administration

Academy of Program/Project & Engineering Leadership

Collaborative Problem-Solving:
The STS-119 Flow Control Valve Issue

Collaborative Problem-Solving: The STS-119 Flow Control Valve Issue

On November 14, 2008, as Space Shuttle *Endeavor* rocketed skyward on STS-126, flight controllers monitoring data during the ascent noted an unexpected hydrogen flow increase from one of the shuttle's main engines. The increased flow did not occur in response to an automated command from the system. Despite this in-flight anomaly, the launch proceeded smoothly—since three flow control valves (one per main engine) work in concert to maintain proper pressure in the hydrogen tank, one of the other valves reduced flow to compensate for the greater flow from the valve that malfunctioned. The likely causes of the malfunction were either an electrical failure or a mechanical failure, which might have resulted from a broken valve. This would require immediate attention as soon as STS-126 landed safely.

The challenge this problem would pose was a familiar one. To ensure the safety of future shuttle missions, management, along with the technical community, would need the best possible analysis to understand the causes of the failure and its implications for future missions. They would have to promote and ensure open communication among the multiple organizations involved in the shuttle program so that all relevant information would be available to decision makers with the responsibility to approve or delay future shuttle flights.

First Impressions

"We knew at least on paper the consequences could be really, really bad, and this could have significant implications for the Orbiter fleet and most urgently the next vehicle in line. Depending on where the vehicle landed, we wanted to get these inspections done and some x-rays done as quickly as we could," said John McManamen, Chief Engineer of the Space Shuttle Program.

The shuttle touched down at Edwards Air Force Base on November 30 after unfavorable weather conditions at Kennedy Space Center (KSC) led flight controllers to divert the landing to California. This delayed work until December 12, when the shuttle was ferried back to KSC aboard a specially equipped 747.

A Complex Assembly and Resupply Schedule

The shuttle program schedule had ten missions left before the planned end of the program in September 2010. These ten missions, which were tightly integrated with international ISS resupply launches by the Progress, ATV, and HTV vehicles and with crew exchange via Soyuz launches, represented a very complex manifest for the Shuttle and Station Programs. The next launch, STS-119, was scheduled for February 12, 2009. Its mission was to deliver to the final set of solar arrays needed to complete the International Space Station's (ISS) electricity-generating solar panels, and to enable the ISS to support its expanded crew of six.

Shuttle and ISS program managers preferred STS-119 to launch prior to mid March so that it would not interfere with the March 26 mission of the Russian Soyuz to transport the Expedition 19 crew to the ISS. If the launch were delayed until after the Soyuz flight, interdependencies in the schedule would require a re-evaluation of other future launches.

Broken Valve

Since the flow control valves are part of the space shuttle Orbiter, discussions began at once between the Orbiter project team and the Johnson Space Center (JSC) and Marshall Space Flight Center (MSFC) engineering organizations about whether to remove the valves from the Orbiter for inspection or to take an x-ray of them in place within the vehicle. The flow control valves involved organizations at JSC, home of the Orbiter Project and the JSC engineering organization, and MSFC, home of the Space Shuttle Main Engine and the Main Propulsion System engineering organizations. Dan Dumbacher, Director of Engineering at MSFC, said, “The flow control valve is in that interesting world of complex interfaces. It can influence what happens on the propulsion side of the equation, but yet it’s owned and the responsibility for the hardware is all on the Orbiter side. Immediately you realize that you’re going to have some complex interfaces between centers, between contractors, and all of the above.” The cultural differences within these organizations helped shape their respective approaches to problem solving, which led to occasional differences of opinion about the best path forward.

An x-ray taken December 19 showed evidence of a problem with a poppet in the valve. Engineers removed the valve and shipped it to VACCO, the only vendor certified to disassemble it. VACCO shipped the disassembled valve to Boeing’s Huntington Beach facility, where inspection determined that a fragment had broken off the poppet, the first time such a problem had occurred during flight. In 27 years, the shuttle program had never experienced a valve failure like this. There had been two similar failures in the early 1990s during testing of a new set of flow control valves for *Endeavour*, but the hardware had always performed as expected in flight. “We knew we had a pretty significant problem well outside our experience base at that point,” said Orbiter Project Manager Steve Stich.

There were a total of twelve flight-certified valves in existence: three in each shuttle, and three spares. Simply buying more was not an option—these custom parts had not been manufactured in years, and NASA had shut down its flow control valve acceptance testing capability at the White Sands Test Facility and at VACCO over a decade earlier.

The Hydrogen Flow Control Valves

During ascent, the external tank provides the main engines with liquid oxygen and hydrogen propellant. As the engines consume the propellant, ullage volume forms (i.e., empty space opens up) within the external tank. To maintain the integrity of the tank, a small amount of the liquid hydrogen is tapped off of each engine's low pressure fuel turbo pump (LPFTP) in the gaseous state and is pumped back into the tank's ullage to maintain the correct pressure. The three flow control valves (one for each engine) regulate this flow of hydrogen gas moving from the engines to the tank.

Each valve has a poppet that resembles a tiny sprinkler head, which pops up and down regulating the flow of gas through the valve. The valves have two rates: high and low. During a launch they switch between the two rates approximately fifteen times.

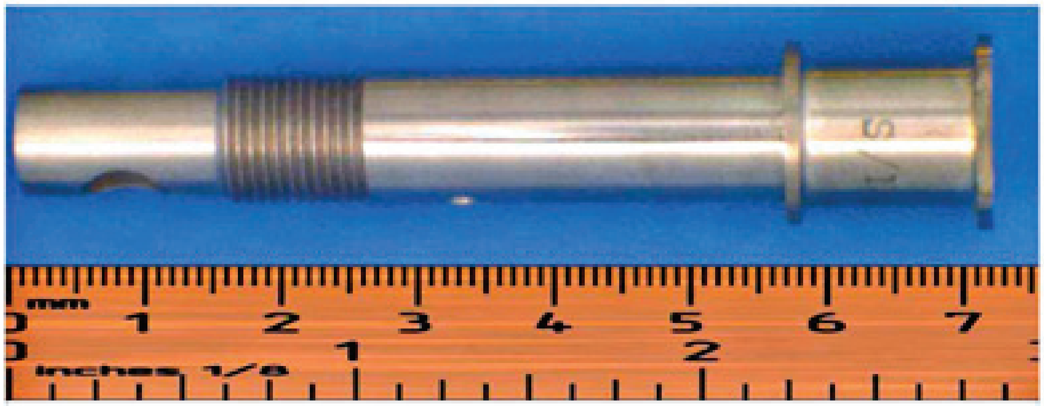

Figure 1. GH2 flow control valve and poppet. (Source: NASA)

Struggling to Bound the Problem

Analysis of the cracked valve showed that the failure resulted from high-cycle fatigue (a condition in which a material sustains damage after experiencing numerous cycles of stress). This raised the question if STS-126 had presented an unusual environment, or if another valve was likely to break in flight due to fatigue-related damage. It also led to complex questions about likelihood and consequences. What was the likelihood of another broken valve? What would be the worst-case consequences of a break? Engineers needed to determine the probable size and the maximum size of a loose particle, understand how it would move through the propulsion system, and what the system could tolerate without experiencing a potentially catastrophic rupture in its lines.

Searching for Cracks

Several teams worked on the problem from multiple angles, including materials, structural dynamics, computational fluid dynamics (CFD), and fracture mechanics. The initial efforts relied on visual inspection and a number of non-destructive evaluation (NDE) techniques, including dye penetrant inspection, magnetic particle inspection, and scanning electron microscope. Electron microscopy could only discover small cracks after the poppet was polished, however, and that polishing invalidated the flight certification of the hardware even as it provided greater ability to inspect it. "A polished poppet could upset the flow balance of the valve rendering it unusable for flow management. In this case the valve could get stuck in the high- or low-flow positions, which could cause a serious issue in flight," said Steve Stich. "In order to ensure that a polished poppet was properly balanced required testing on the flow test system that had been shut down at the White Sand Test Facility in the mid-1990s. So we were in a bit of Catch-22 situation with respect to performing the best possible NDE."

This meant it could only be used on non-flight hardware, such as an older configuration of the valve known as a -361 valve. Initially the scanning electron microscope inspections discovered no cracks in all eleven unpolished flight poppets and in all but one of the older -0361 poppets. "We felt like that was pretty good flight rationale in that the problem was not pervasive throughout the fleet" said Steve Stich. Around the same time, though, testing of a polished -361 valve did reveal cracks, raising questions about the value of the inspection technique for valves that would actually fly.

One NDE technique that was initially dismissed was the use of an eddy current system because the size of the probe head was too large for the hardware.

February 3, 2009: Flight Readiness Review #1

With the launch scheduled for February 19, the program scheduled a Flight Readiness Review for February 3. At the review, it quickly became clear that the engineering and safety organizations felt that significant work needed to be done before sound flight rationale could be established. Steve Altemus, Director of Engineering at Johnson Space Center, summarized the knowledge gap from the JSC engineering community's point of view: "We showed up at the first FRR and we're saying, 'We don't have a clear understanding of the flow environment, so therefore we can't tell you what the likelihood of having this poppet piece come off will be. We have to get a better handle on the consequences of a particle release.'" The most important outcome of the meeting was the establishment of new lines of inquiry that could lead to better understanding.

"At the end of that FRR, we laid out some criteria that I thought were important to understand before we went to 'go fly,' and left that with teams to go work," said Bill Gerstenmaier, who chaired the FRR in his role as Associate Administrator for Space Operations.

On February 6, the launch was delayed until February 22.

The Flight Readiness Review

As described in NASA Policy Directive 8610.24B, a Flight Readiness Review "is held to update the mission status, close out actions from the previously held LVRR [Launch Vehicle Readiness Review] and Customer MRR [Mission Readiness Review], and certify the readiness to proceed with initiation of the launch countdown." As the definition suggests, the FRR typically evaluates work done on issues identified at earlier reviews, and gives the teams responsible for various aspects of the mission an opportunity to make sure those technical questions have been adequately dealt with and to raise any additional concerns.

Impact Testing

After the first FRR, the Orbiter Project authorized three different lines of impact testing to learn more about whether a particle would puncture the pressurization lines downstream from the valve. At Glenn Research Center, testing focused on the material properties of the flight hardware and the impact of particles striking the material at a certain velocity and orientation, along with testing several Orbiter and ET components from the same material as in the actual gaseous hydrogen flight hardware. JSC Propulsion Systems Branch Chief Gene Grush traveled to Stennis Space Center to set up a test stand that fired particles through a full-scale mock-up of the propulsion system within the External Tank. Grush also coordinated efforts with White Sands Test Facility to run a similar test stand focusing on the Orbiter part of the system near the flow control valve exit.

The Orbiter analysis divided the Orbiter's Main Propulsion System into seventeen discrete areas that a particle would travel through and considered the type of material used and the thickness of the wall in each area. The data from these tests and other analyses contributed to a probabilistic risk assessment of the entire flow control valve hydrogen repress system.

At the same time, the computational fluid dynamics (CFD) analysts improved their characterizations of the environment inside the propulsion line, figuring out the velocity and spin of a given sized particle as well as the probable path it would travel through the elbow-joint turns in the pipe.

As data began to come in from these tests, the program decided to convene a second FRR on February 20, 2009. Prior to the meeting, some members of the engineering and safety organizations expressed doubts about the timing of the review. Chris Singer, Deputy Director of Engineering at Marshall Space Flight Center, told Steve Stich that the charts depicting the Orbiter Project's flight rationale did not bode well for the review.

Scott Johnson, Chief Safety and Mission Assurance Officer for the Shuttle, thought the review was premature. "The majority of the safety community was concerned about the amount of open work in front of us, and as a result I recommended that we delay the FRR," he said. "We were doing impact testing that was not complete and was not due to be complete until basically the day of the FRR. We still had a lot of the analysis work going on. We weren't really that close to being able to quantify the risks."

Senior leaders of the shuttle program and engineering teams had a meeting prior to the review with FRR board about what to expect. "We talked a lot about flow control valves," said Steve Altemus, "but what we didn't talk about were the agency-level constraints, and leveling risk across the programs and the international partners—what it really meant to delay flight past the Soyuz cutout versus taking this risk of flying with the potential to break this poppet."

Right before the second FRR, the engineering team learned about deficiencies in the processes used to inspect the poppets. Joyce Seriale-Grush, Orbiter Chief Engineer, said, "Before we went into the FRR, a few of us were pretty much of the opinion that we should stand down and take these valves off, refine our inspection processes further, and either put ones on that had been inspected with the new processes or re-inspect these."

February 20, 2009: Flight Readiness Review #2

The second FRR for STS-119, on February 20, 2009, was far from typical. The session lasted nearly fourteen long hours, and the outcome was not clear until the end. "It was much more of a technical review than typical Flight Readiness Reviews. There was lots of new data placed on the table that hadn't been fully vetted through the entire system. That made for the long meeting," said FRR chairman Bill Gerstenmaier.

Well over one hundred people were in attendance at the Operations Support Building II at Kennedy Space Center, seated around the room in groups with their respective organizations as technical teams made presentations to the senior leaders on the Flight Readiness Review board. Some participants believed that the analysis done on the potential risk of a valve fragment puncturing the tubing that flowed hydrogen to the external tank to the shuttle main engines showed that the risk was low enough to justify a decision to fly. Others remained concerned throughout that long day about the fidelity of the data, and that they didn't know enough about the causes of the valve failure and the likelihood and risk of the failure occurring again.

Despite the tremendous amount of analysis and testing that had been done, technical presentations on the causes of the broken valve on STS-126 and the likelihood of that happening again were incomplete and inconclusive. Unlike at most FRRs, new data such as loads margins computations that couldn't be completed in advance streamed in during the review and informed the conversation. A chart reporting margins of safety included "TBD" (to be determined) notations. The Orbiter Gaseous hydrogen line summary chart displayed several areas where the worst-case loads margins exceeded the standard 1.4 factor of safety, which allows for material properties variations for components manufactured from different lots of material.

Doubts about some of the test data arose when Gene Grush received a phone call from Stennis Space Center informing him that the test program there had used the wrong material. "I had to stand up in front of that huge room and say, 'Well there's a little problem with our testing. Yes, we did very well, but the hardness of the particle wasn't as hard as it should have been.' That was very critical because that means that your test is no longer conservative. You've got good results, but you didn't test with the right particle," he said.

Ralph Roe, director of the NASA Engineering and Safety Center (NESC) and a longtime veteran of shuttle FRRs, noted how unusual this review proved to be:

> "Usually the projects come to the flight readiness review with all their flight rationale worked out and you hear it and there may be a question or two. But in this particular case they obviously hadn't gotten to complete flight rationale yet so there was pretty enthusiastic debate. Different people had different opinions about what the data meant, and they were able to voice that in that forum, which was good."

NASA Chief Safety and Mission Assurance Officer Bryan O'Connor characterized the openness of the discussion as excellent. "Gerst [FRR chairman Bill Gerstenmaier] was absolutely open. He never tried to shut them [the participants] down. Even though he could probably tell this is going to take a long time, he never let the clock of the day appear to be something that he was worried about. I thought, that's bad in one way, it says we're probably going to have a long day, but it's good in another, and that says that you don't have the chairman of this panel putting undo pressure on people to sit down or be quiet," he said.

Toward the end of the meeting, Gerstenmaier, who had spent the day listening to presentations and eliciting comments, spoke about the risks of *not* approving Discovery's launch: risks to the International Space Station program and to the shuttle schedule. A few participants perceived Gerstenmaier's comments as pressure to approve the flight. Others saw it as appropriate context setting, making clear the broader issues that affect a launch decision. After he spoke, he gave the groups forty minutes to "caucus," to discuss what they had heard during the day and decide on their recommendations. When they came back, he polled the groups: "go" or "no-go."

The Process Works

When Gerstenmaier polled the room after the break, the engineering and safety organizations and some Center Directors in attendance made it clear that they did not find adequate flight rationale. Many felt there were too many uncertainties that the extensive, ongoing testing and analysis had not yet resolved. With no precise way of calculating the level of risk the flight faced, a launch decision could not be justified.

"As a community, we never really got our arms around the true risk," said Steve Altemus. "There were varying degrees of uncertainty in all the different pieces of analysis and test data that were out there, both in the likelihood of occurrence and the consequences of failure."

"One of the key tenets of flight rationale that most individuals and organizations were looking for was a maximum bound on the potential particle release size. This statement was repeated by many of the board members in the final discussion going around the table. This action had been in work for some time, but proved to be a difficult problem to solve. We knew this answer was important and would play hand in hand with the other elements of flight rationale especially in understanding the consequence and risk of release," said John McManamen.

Bill McArthur, Safety and Mission Assurance Manager for the Space Shuttle at the time, said, “The fact that people were willing to stand up and say ‘We just aren’t ready yet,’ is a real testament to the fact that our culture has evolved so that we weren’t overwhelmed with launch fever, and people were willing to tell Bill Gerstenmaier, ‘No, we’re no-go for launch.’”

Several participants thought it might have been better to break off the meeting and reconvene the following morning—it was unquestionably an over-long, exhausting day, and people would undoubtedly have felt sharper after a night’s sleep. But no one thought that the outcome would have been different.

As the participants filed out of the meeting, Joyce Seriale-Grush said to NASA Chief Engineer Mike Ryschkewitsch, “This was really hard and I’m disappointed that we didn’t have the data today, but it feels so much better than it used to feel, because we had to say that we weren’t ready and people listened to us. It didn’t always used to be that way.”

Moving Forward

With the launch postponed after the second FRR, Bill Gerstenmaier had doubts about the likelihood that the work could be completed in time to make the Soyuz cutoff date. His experience told him to reserve judgment. “Rather than me make the random decision to go move somewhat arbitrarily at this point based on the data we saw in the meeting and where I thought we would be, (I decided) I’m going to go ahead and kick it back to the team, give them the action, see what they can go do and see how it comes out,” he said.

The testing and analysis continued on all fronts. In addition to the work at JSC, MSFC, Glenn, and White Sands, there were efforts across the country starting well before the second FRR. The Boeing facility in Huntington Beach, VACCO, the NASA Shuttle Logistics Depot (NSLD) in Florida, Pratt Whitney Rocketdyne, and Ames Research Center were all engaged, and experts from the NASA Engineering and Safety Center (NESC) provided support to the NASA engineering teams. At the peak, roughly one thousand people were working to solve the problem.

A New Inspection Tool

Early in the investigation, the eddy current inspection technique had been ruled out because the probe head was the wrong size for the job and because of magnetism concerns. Charles Bryson, an engineer at MSFC, used his eddy current probe equipment, with a relatively large probe head, to inspect a poppet sent to MSFC for fractography review from the Boeing facility at Huntington Beach. His inspection, as confirmed by fractography work, indicated that the eddy current inspection technique showed promise in finding flaws.

While at the FRR on Friday, February 20, Rene Ortega, Propulsion Systems Engineering and Integration Chief Engineer at MSFC, told colleagues from the Materials and Processes (M&P) Problem Resolution Team (PRT) about Bryson’s inspection results and offered Bryson’s expertise to inspect poppets at Boeing’s facility. After a few phone calls, Ortega helped arrange for Bryson to examine several poppets at the Huntington Beach facility with his eddy current setup if he could fly out over the weekend. Bryson traveled and remained there refining the technique for several days. He then worked collaboratively with a team from JSC led by Ajay Koshti, a non-destructive evaluation (NDE) specialist with expertise in eddy current investigations. Koshti brought an eddy current setup with a better response than Bryson's, and together using the more suitable setup, they arrived at a consistent inspection technique.

Ortega explained how the new eddy current technique supplied a missing piece of the puzzle. "Once we were able to screen flaws with the eddy current and there wasn't a need to polish poppets with that process, then we had a method by which we could say that we had a certain size indication that we thought we're pretty good at screening for non-polished poppets."

Through fracture analysis, engineers had found that some of the smaller flaws identified in the poppets didn't seem to be growing very fast. "Through that exercise, we came up with the suggestion that, 'Hey, it doesn't look like these flaws are growing out very rapidly in the flight program, and with the screening of the eddy current we can probably arrive at a flight rationale that would seem to indicate that those flaws being screened by the eddy current wouldn't grow to failure in one flight,'" he said. In short, the eddy current technique was not a silver bullet, but in conjunction with the other techniques and test data, it provided critical information that would form the basis for sound flight rationale.

Steve Altemus thought Koshti's efforts exemplified the engineering curiosity that NASA needs to succeed, and he viewed it as an engineering leadership responsibility to create an environment in which new and diverse ideas could receive a fair hearing. "It's important to recognize that we're not always the smartest one in the room, that perhaps there's somebody over there in the corner of the room, and that we have to pull out of them what their thoughts are, because they've got the answer. He had the answer."

March 6, 2009: Flight Readiness Review #3

With a full complement of analyses and the results from the test programs all supporting a shared understanding of the technical problem, there was wide consensus among the community that the third Flight Readiness Review would result in a "go" vote.

"By the time we eventually all got together on the last FRR the comfort level was very high," said Bryan O'Connor. "For one thing, everybody understood this topic so well. You couldn't say, 'I'm uncomfortable because I don't understand.' We had a great deal of understanding of not only what we knew about, but also what we didn't know about. We had a good understanding of the limits of our knowledge as much as possible, whereas before we didn't know what those were."

Steve Stich summarized the progress that had been made. "By the third FRR, there was no new test data or analysis coming in late. We had better characterized the risk of damage in the Orbiter and ET due to a poppet fragment through our impact testing and stress analysis. We had better characterized the worst case even if the poppet fragment ruptured the line, along with what hole sizes would be required to cause enough hydrogen to meet the flammability limits in the ET and in the orbiter. Overall, we had a much better characterization of the risk by the time we got to that third and final FRR. Plus, we were able to use this new eddy current technique to say with more certainty that the poppets did not have any significant cracks prior to launch. Even though we didn't have total root cause, we knew we weren't starting with large cracks. We had also begun some materials testing at Marshall that showed it was very unlikely to grow a crack from a very small size to failure in one flight. So we ended up having extremely good rationale for that third FRR."

The Flight Readiness Review Board agreed, and at the third and final FRR, STS-119 was approved to launch on March 11, 2009.

Epilogue

After delays due to an unrelated leak in a liquid hydrogen vent line between the shuttle and the external tank, STS-119 lifted off on March 15, 2009.

Two months after the completion of the mission, Bill Gerstenmaier spoke to students at Massachusetts Institute of Technology (MIT) about the flow control valve issue. In an email to Shuttle team members, he shared a video of the lecture and wrote, "I am in continue to learn mode. There is always room to improve."

Appendix 1. Timeline

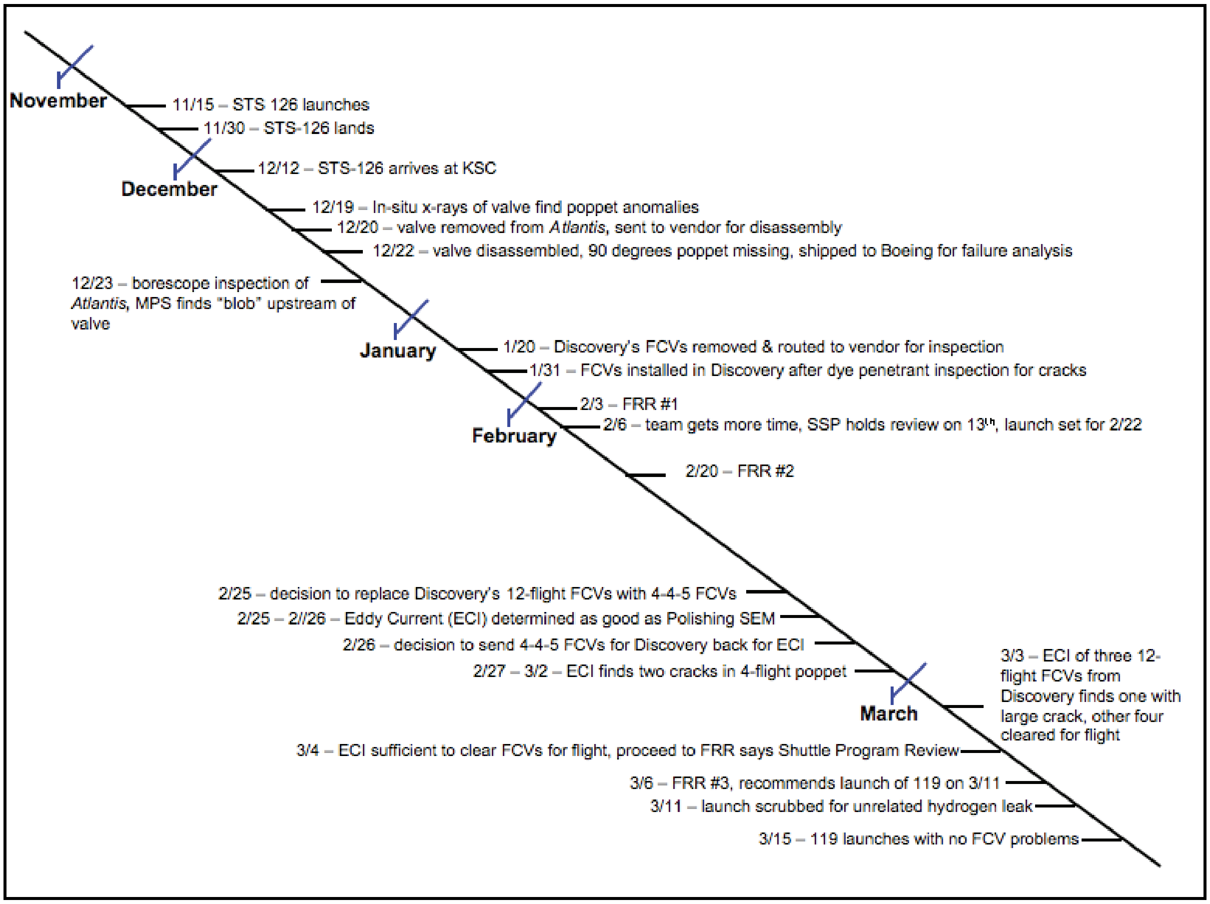

Figure 2. Timeline of key events and developments.

Appendix 2. Reflections on Technical Leadership

Leaders took different lessons from the experience of resolving the flow control valve issue prior to STS-119.

“I had better be empowering and I had better be motivating folks to stay hungry…to keep digging, keep looking, and figure out what the heck is really going on if I want a good decision.”

- Bill Gerstenmaier, Associate Administrator for Space Operations

“As a leader, it’s my responsibility to create an environment that is open to ideas and diverse perspectives, to say, ‘Hey, I’m not necessarily the smartest guy, show me what you know,’ and then be able to act on it and get an audience so that the broader community understand that there’s a good solution out there.”

- Steve Altemus, Director of Engineering, Johnson Space Center

“The one thing I would take away from this effort is that NASA can solve any problem if the right experts are involved and that Centers with different perspectives and cultures can work extremely well together. To have developed flight rationale for this problem with only a one month delay to the STS-119 mission, while preserving the rest of the manifest was truly an amazing feat by this incredibly talented team.

In retrospect, as the Orbiter Project Manager, I should have enlisted the assistance of the materials experts from the MSFC much earlier in the process, since ultimately they had the proper expertise to deal with high cycle fatigue. Additionally, we should have established a better mechanism for establishing communication of fast developing information from the beginning of this failure investigation so that the entire team could remain equally informed.”

- Steve Stich, Orbiter Project Manager

“Learning how to create some alternate trade space for this very disciplined engineering rigorous community is an art, but it’s definitely worth doing. We fly complex systems and the rigor usually works, but sometimes we have to have a little different trade space. It takes some special people who have some multi-discipline experience to be able to start creating that trade space.”

- Chris Singer, Deputy Director of Engineering, Marshall Space Flight Center

Teaching Notes

This case study has been designed for use in a classroom setting. Please read the full case prior to in-class discussion to allow ample time for analysis and reflection.

Consider the following questions:

- What role did cognitive diversity—the extent to which the group reflected differences in knowledge, experience, and perspectives—play in the resolution of the technical problem?

- How did different participants in the case exercise leadership at different points in time?

- How did communications among stakeholders help shape the eventual outcome of the case?

Ask participants to discuss in small groups, encouraging them to draw analogies to their own experience and develop as many interpretations as possible. The small groups will then reconvene as a large group and share their conclusions.