the Availability Digest

Triple Redundancy Failure on the Space Station

November 2007

NASA photo gallery

In June, 2007, a triply-redundant attitude and environmental control computer provided by Russia failed on the International Space Station (ISS). Had this been a mission to Mars, it would have been fatal. Only the space station's proximity to Earth, which put it in range of support and resupply missions, prevented a tragedy.

Though the problem was circumvented in a few days by the space station crew, it took weeks for the station crew and ground engineers to determine the source of the problem. It turned out to be caused by a single point of failure in the otherwise triply-redundant system – a failure that was highly unlikely but one that occurred anyway.

During this experience, many technological and diplomatic lessons were learned. Understanding the interaction between the crew members was imperative since the U.S. and Russia are bound to be partners in a Mars mission.

The Shuttle Mission

The U.S. space shuttle Atlantis docked with the ISS on June 10, 2007.[1] Commanded by Colonel Frederick Sturckow, it carried the second and third starboard trusses required to extend the space station as well as an additional pair of solar arrays to provide additional power for the station. It also carried U.S. astronaut Clayton Anderson who would stay on the ISS and replace Sunita Williams, who would return to Earth.

During the mission, the trusses were installed on the space station along with the new solar arrays. The solar arrays were unfurled and their power added to the space station power.

From the left (front row) are astronauts Rick Sturckow, commander, Lee Archambault, pilot, and Patrick Forrester, mission specialist. From the left (back row) are astronauts John Olivas, Clayton Anderson, Jim Reilly, and Steven Swanson, all mission specialists. Image credit: NASA

Disaster Strikes

During Atlantis' visit to the space station, and after

[1] NASA mission report, http://www.nasa.gov/mission_pages/shuttle/shuttlemissions/sts117/main/index.html

the new trusses had been installed and the deployment of the new solar arrays had been completed, a Russian triply-redundant computing system that controlled space station attitude and environment suddenly failed.[2] This system was critical in holding the space station stable during the undocking of the shuttle. Without this system, the shuttle might not be able to leave the space station. In addition, the space station's atmospheric control system was disabled.

It was quickly determined that the failure was caused by the simultanous loss of power to all three computers. Power had been shut off by a surge protection unit designed to protect the computers from power surges beyond the capabilities of their own power filters. A NASA internal technical report describing this failure said, "On 13 June, a complete shutdown of secondary power to all (three) central computer and terminal computer channels occurred, resulting in the loss of capability to control ISS Russian segment systems."

Finger Pointing

Russian officials were quick to blame NASA for "zapping their computers" with "dirty" 28-volt power from the newly-installed solar array. Another theory was that the expanded station structure following the truss installations might be excessively charging up due to its orbital speed through Earth's magnetic field. These were the first of many bad guesses by top Russian program managers and would distract engineers trying to get to the real source of the problem.

A Jury-Rigged Fix

In the meantime, the computers had to be fixed – and fast. The station crew assumed that some external interference such as noise in the 28-volt power supply was responsible for generating false commands inside the computers' power-monitoring system and caused it to send shut-down commands to all three computers. Based on this reasoning, they bypassed the power monitoring system to two of the computers by using jumper cables. These two computers were now subject to damage by power surges, but by now the power system had settled into a steady state and was generating clean power.

By the time the shuttle undocked on June 19, the computers seemed to be operating properly; and undocking proceeded without incident. Replacement parts were quickly sent via a robot supply ship while analysis teams on the ground wrestled with why the system failed and why the jumper cables corrected the problem. Was the problem really fixed, or could it happen again?

Root-Cause Analysis

The ground engineers were not the only ones working on the problem. On board the space station, the Russian commander Fyodor Yurchikhim and astronauts Oleg Kotov and Clayton Anderson spent their time disassembling the power control boxes and the associated cabling in order to look for clues that might lead to the cause of failure. Though multiple scopes and probes failed to find the problem, their eyes and fingers did.

What they discovered was that the connection pins from the power monitoring unit were wet and corroded. Continuity checks showed that the command lines in the cable coming off the unit had failed. Even worse, one of the command lines had shorted. It was the power-off command line that went to all three computers. The shorted condition created the disastrous power-off command. The jumper cables had bypassed the false power-off command and had allowed the computers to function properly once again.

---

[2] James Oberg, Space Station: Internal NASA Reports Explain Origins of June Computer Crisis, IEEE Spectrum Online; October, 2007 (http://www.spectrum.ieee.org/oct07/5598).

But what had caused the corrosion? Water condensation, one of the most common problems in space avionics. The problem was traced to a malfunctioning dehumidifier. As NASA later reported, the problem presumably was "the result of repeated emissions of condensate from the air separation lines" of a nearby dehumidifier. Though the power usage of the cables was supposed to keep them warm enough to prevent condensation, the malfunctioning dehumidifier's frequent on-off cycles led to surges of water vapor. The situation was aggravated by a stream of cold air from another location on the dehumidifier that at times cooled the cables below the dew point at which moisture could condense.

Chewing Gum and Duct Tape

Replacement units for the faulty components were quickly built by the European Space Agency and sent to the space station on a robot supply ship.

In August of 2007, the space shuttle Endeavor docked at the space station. The station's stabilization function was turned over to Endeavor, and the old computer network was torn down. The crew found condensate mixed with mold behind the computer units.

While installing the new units and their cables, it was found that one cable was 40 centimeters too short. After careful inspection, the crew could find no evidence of corrosion on the old cable, so they decided to continue its use.

As further protection, the crew rigged a thermal barrier between the computers and the dehumidifier. The thermal barrier was built using a surplus reference manual and ordinary gray tape.

Finally, they removed the jumper cables, verified that the system was working, and closed the access panels. Mission accomplished.

The Fatal Design Flaw

Once the problem was understood, it became clear that the system suffered from a fatal design flaw. The supposedly triply-redundant design included a single point of failure – the external power monitoring unit that, by itself, could turn off all of the computers. Should it fail (as it did due to condensation), the triply-redundant system was down.

A proper triply-redundant system design would make sure that no single failure (or even a dual failure) could take down the system. In this case, this condition could have been satisfied by providing a separate power monitoring unit for each computer.

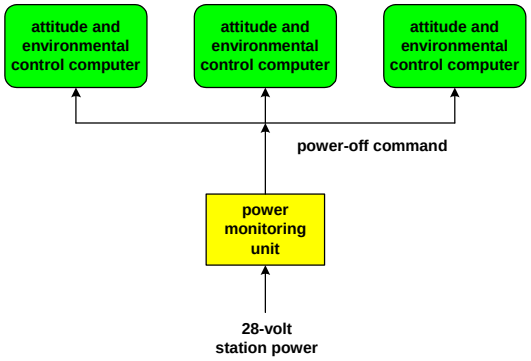

Single Point of Failure

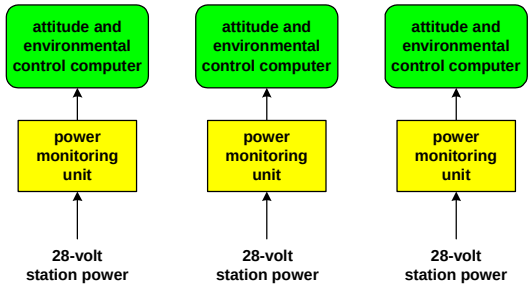

Proper Triply-Redundant Design

Lessons Learned

There are two main lessons to be learned from this experience:

- One is the obvious. If you are designing a redundant system, make sure that there are no single points of failure. It is inconceivable that any good design engineer would purposefully make such an error. However, single points of failure can be insidious; and they may never be discovered despite exhaustive testing. It is only through careful peer design review that such design faults can be uncovered.

- The other lesson is that finger-pointing should never be allowed. “This cannot possibly be my fault. It has to be yours” is an attitude that is unacceptable. Until the problem is resolved, it is everyone’s problem. Everyone must assume that the problem lies within his area of responsibility and must continue to search for the cause until it is found. In this case, the obvious cause was a malfunctioning power source installed by NASA. However, it turned out that the problem was an improper design on the part of Russian engineers, one that was aggravated by a faulty dehumidifier.

These lessons ignore the faulty design and placement of the dehumidifier. Though it ended up being the ultimate culprit, a proper redundant design would have been tolerant of this malfunction.