PROMOTING RESILIENCE IN HUMAN SPACE FLIGHT AT A TIME OF FISCAL PRESSURE

C.W. Johnson [(1)], Michael P. Fodroci [(2)], A. Herd [(3)], M. Wolff [(4)],

[(1)] Department of Computing Science, University of Glasgow, Scotland.
http://www.dcs.gla.ac.uk/~johnson, Email: Johnson@dcs.gla.ac.uk
+44 141 330 6053 (Tel.), +44 141 330 4913 (Fax).

(2) NASA Johnson Space Centre, International Space Station Division.
Safety and Mission Assurance, Mail Code NE, Houston, USA.
Email: michael.fodroci-1@nasa.gov
+1 281 483-4206 (Tel.), +1 281 660-0299 (Mobile)

[(3)]ESA Operations Safety Unit, D/OPS-H & ESA Independent Safety Office, ESTEC,
Keplerlaan 1, 2201 AZ Noordwijk, The Netherlands.
Fax: +31 71 565 6278, Phone: +31 71 565 6745, Mobile: + 31 650 685425

[(4)] Software Systems Division, Directorate of Technical and Quality Management,
ESTEC Keplerlaan 1, 2201 AZ Noordwijk, The Netherlands.
Fax: +31 71 565 5420, Phone: +31 71 565 3206

ABSTRACT

Space missions require significant investments to develop and sustain the underlying engineering infrastructures. Assuring mission success (Return-On-Investment) also depends upon investments in the training that supports closer integration between flight crews and ground teams. However, economic and fiscal pressures are forcing many governments to demand savings from their national space programs. From an engineering perspective, this makes it essential to identify those areas of investment that contribute most to the resilience of space missions. The following pages analyze a number of successful interventions by flight crews and ground teams to resolve problems, including but not limited to hardware failures, on the International Space Station. These case studies are used to identify ways in which finite investments might best be deployed to promote the resilience of future missions.

1. INTRODUCTION

Many governments face significant fiscal pressures that are placing constraints on their civil space agencies. This has led to the cancellation or curtailment of long term programs, including NASA's Constellation initiative. ESA's spending in 2010 and 2011 has been frozen at approximately €3.7bn. Some member states, including Ireland, Portugal and Spain, have experienced considerable difficulties in securing their individual subscriptions. In the United States, the budget deficit has fuelled Republican hostility to the administration's plans for the integration between Federal and commercial space programs. It seems likely that elements in Congress will try to cut subsidies for commercial human space flight from $6 billion over six years to $3 billion. It is against this background that a joint project was developed between ESA, NASA, the US Air Force and the UK Engineering and Physical Science Research Council to identify techniques that support the resilience of space-based operations at a time of financial stringency (Johnson, Herd and Wolff, 2010, Johnson, Fletcher, Holloway and Shea, 2009). The aim of assuring resilience of space-based operations is to promote those behaviors that promote mission success (including safety of crew and vehicle). In particular, this paper focuses on the ways in which pre-mission planning and flight and ground team training support the flexible interventions that characterize improvised responses to complex systems failures. In this way even when the unexpected occurs on-orbit operations may continue without impacting overall achievement of mission success, and assure on-orbit resources are not over-assigned to addressing unplanned events.

A Brief Overview of Pre-Flight Training: It is impossible to provide a complete account of the pre-mission preparations that support human space flight. The following paragraphs, therefore, provide a very high-level summary of ISS training from a US perspective. The intention is to provide an impression of the scale of investment required to support the pre-flight phases of human space flight. The remaining sections of this paper use a number of key incidents on the International Space Station (ISS) to provide some indication of the Return on Investment when resilience techniques are applied.

The ISS crew requires a minimum of 18 months training prior to a mission. The precise time depends on

whether or not the individual has specific language skills, such as fluency in Russian as well as English, or whether they have acted as backup for a previous crew member. A Crew Qualifications and Responsibility Matrix is, typically, developed once a crew member has been assigned to a flight. This provides high level details about the tasks that they will be expected to perform. It also serves as an enumeration of the skills that they will have to demonstrate before launch. This matrix distinguishes between operators and specialists. All crewmembers must be qualified to operate all of the main ISS infrastructures. Operators are expected to use particular systems. In contrast, specialists have additional training. They must be able to understand sufficient details of the component architecture to be able to diagnose and respond to a range of potential failures.

A training team is assigned to the crew and together they devise the program that is intended to provide the skill sets that are identified in the Crew Qualifications and Responsibility Matrix. There are instructors for each of the main ISS infrastructures and additional teams for the scientific experiments, for the operation of the robotic arm, for medical interventions and for Extravehicular Activities (EVAs). The costs associated with EVA training are significant. Crews must learn a range of theoretical and practical skills both in order to conduct the activity and also to support their colleagues when they are outside the ISS. Neutral buoyancy tanks that include scale models of the ISS and the Orbiter payload bayare use d to provide individuals with an idea of what it would be like to work in the suits, at the Gagarin Cosmonaut Training Center in Star City and at the Johnson Space Center in Houston, Texas. These exercises provide ground teams to assess the physiological characteristics of crewmembers; different individuals will use their oxygen supply at different rates even though they perform similar tasks. These exercises are also used to assess cognitive resilience and cooperation between teams in response to system failures. In the past, around seven hours of training have been provided in a neutral buoyancy tank to rehearse every hour of operations in an eventual EVA. For some individuals this equates to more than 100 hours of training in facilities

As might be expected for an international mission, portions of the training take place at facilities in several countries based primarily on where the technical expertise for that training resides. For example, there are specialist facilities for working with the robotic arm in Canada. Periods in Russia provide the crew with experience of working in a foreign language. It also provides first hand opportunities to talk with the specialist engineering teams who maintain key infrastructure components including the Environmental Control and Life Support Systems and the Russian Command and Data Handling computers; both of which will figure prominently in the following case studies of critical incidents.

During the training, emphasis is placed on identifying and implementing the Standard Operating Procedures (SOPs) that guide everyday interaction on the ISS. One aspect of this is the use of the Inventory Management System that helps crew members to identify necessary equipment and supplies. The crew also receives detailed training in the operation and/or maintenance of the U.S. Command and Data Handling systems, the Electrical Power System as well as the Mobile Transporter that is used to move the ISS robotic arm. Further training focuses on the Caution and Warning systems – a system that provides the crew with a visual and audible indication that they are required to take action; which again will figure in the case studies that form the remainder of this paper. A further element of the training that continues throughout all of the exercises is the interaction and coordination of flight crew operations and ground support teams. This is critical because the ISS crew and Flight Director (who leads the mission) can call upon a vast array of engineering and other technical expertise. The typical Mission Control Centre flight control team positions for the space station include:

- Assembly and Checkout Officer (ACO)
- Attitude Determination and Control Officer (ADCO)
- Communication and Tracking Officer (CATO)
- Environmental Control and Life Support System (ECLSS)
- Extravehicular Activity Officer (EVA)
- Flight Director
- Flight Surgeon
- Integration Systems Engineer (ISE)
- Onboard, Data, Interfaces and Networks (ODIN)
- Operations Planner (OPSPLAN)
- Operations Support Officer (OSO)
- Power, Heating, Articulation, Lighting Control Officer (PHALCON)
- Remote Interface Officer (RIO)
- Robotics Operations Systems Officer (ROBO)
- Thermal Operations and Resources (THOR)
- Trajectory Operations Officer (TOPO)
- Visiting Vehicle Officer (VVO)

The pre-mission simulations help crew members to draw upon the expertise from each of these individuals and their support teams. Mission control must, in turn, practice with the crews to correctly identify how best to support their needs during a range of critical scenarios.

In many cases, the individuals listed above must coordinate their interventions with a range of other teams – for example between Russian and US mission control. As the crew gets closer to their flight, they begin to train with the US Space Shuttle (Orbiter) teams that will be responsible for taking them to the ISS. It is important to stress that this phase can introduce different structures and responsibilities for individuals. The precise details depend on whether the ISS crewmember will form part of a Shuttle or a Soyuz mission. However, some of the differences in terms of the allocation of tasks can be illustrated by comparing the ISS mission control responsibilities with those of the Shuttle Flight Control Positions:

- Assembly and Checkout Officer (ACO)
- Booster Systems Engineer (BOOSTER)
- Data Processing System Engineer (DPS)
- Emergency, Environmental, and Consumables Management (EECOM)
- Electrical Generation and Integrated Lighting Systems Engineer (EGIL)
- Extravehicular Activity Officer (EVA)
- Flight Activities Officer (FAO)
- Flight Dynamics Officer (FDO or FIDO)
- Ground Controller (GC)
- Guidance, Navigation, and Controls Systems Engineer (GNC)
- Instrumentation and Communications Officer (INCO)
- Mechanical, Maintenance, Arm, and Crew Systems (MMACS)
- Payload Deployment and Retrieval System (PDRS)
- Propulsion Engineer (PROP)
- Rendezvous (RNDZ)
- Trajectory Officer (TRAJ)
- Transoceanic Abort Landing Communicator (TALCOM)

Later sections will describe how the allocation of tasks and responsibilities between these positions and between ground support and the flight crew have a profound impact on the resolution of systems failures. When time is limited and the detailed causes of a warning cannot accurately be identified, it is critical that each member of each team works to avoid the omission of necessary tasks and the unnecessary duplication of essential operations.

Impact of Financial Pressures on Pre-Flight Training: Many International Partners are facing fiscal constraints that, in turn, have significant effects on their ability to resource their civil space programs. This has knock-on effects for international programs when, for example, a cutback in one state will affect the training that they can provide to the crews and ground teams for other project partners. Other problems are created in ensuring an adequate distribution of funds across complex international space programs where, for instance, states that are suffering fiscal pressures may also be called upon to take a greater role in the technical support of future missions. Evidence for these assertions can be seen in some of the complexities that have arisen towards the end of the Shuttle program.

All partners in the ISS recognize the need to maximize the returns for public investments in human space flight. However, the search for fiscal efficiency can also bring with it organizational changes that have significant engineering implications. This can be seen in the cancellation of Constellation and the promotion of commercial space flight; including the outsourcing of crew transportation to the International Space Station. The promotion of external contracting for crew transportation to and from Earth orbit did not originate with the Obama administration. George W. Bush's Commercial Orbital Transportation Systems created a framework in which initial funding would be increased as particular milestones were met. This program included provision firstly for the delivery of cargo and then of people to the ISS. This initiative became increasingly important with the cancellation of Constellation, including the Ares 1 system that was the only alternative to these commercial ventures. As we write this paper, attention is focused on a small range of commercial space companies including SpaceX, Orbital Systems, and Boeing. In the meantime, many of the individuals and teams with significant expertise in training for space missions have found alternate employment in a period of considerable uncertainty.

Similar concerns have affected the Russian space program. There has been a gradual transfer of responsibility for critical infrastructure, including the Gagarin Cosmonaut Training Center in Star City, from the Russian Ministry of Defence to the Roskosmos civilian space agency. This transfer created shortages in some training roles because key individuals preferred to continue their careers within the military. It is hard to underestimate the impact that such disruptions have on the engineering and management of human space flight given the specialized and skilled nature of these operations. It is simply not possible to issue a conventional job advert and expect a crop of well qualified applicants, when the required competencies take decades to acquire.

Organizational change also creates uncertainty. This has undermined attempts to retain key personnel/competence during the interregnum between military and civil control or, in the US case, between Federal and commercial space operations. In the meantime, there is a continuing need to fulfill the

training requirements within the Crew Qualifications and Responsibility Matrices, mentioned in previous sections.

Organizational change and financial uncertainty not only affects the staffing of training centers and engineering teams. It also affects the physical infrastructures that are available to support pre-flight preparations. As the Shuttle program nears its end a host of training and simulation facilities have either been mothballed or are in the process of being dismantled. Some of these capabilities will certainly be required again when commercial missions are scheduled for the ISS. There are further similarities with the Russian experience at Star City where many facilities were starved of cash in the anticipation that civil operations might bring commercial funding to supplement existing resources. This has occurred at a time when Russians have had to increase their training provision to support additional missions with the retirement of the Shuttle.

Fiscal constraints and the associated organizational changes have created operational concerns. For instance, in the past ISS crews have been trained to operate both the US and Russian EVA systems, including spacesuits and airlocks. However, it has been difficult to retain this practice in the face of funding cuts on the training infrastructures. It is possible that in the future this will reduce the redundancy that has protected ISS operations. There are also concerns about the quality of training that some crews are receiving. ISS partners have, therefore, continued to monitor crew performance in pre-flight certification tasks compared to previous generations of ISS crews.

2. ISS URINE REPROCESSING ASSEMBLY CASE STUDY

The following pages consider the many different ways that flight crews and ground teams cooperate to resolve degraded modes of operation. This analysis is intended to help focus finite training resources on behaviors that have promoted safe and successful operations. We are also concerned to protect the budgets available for the pre-mission phases (in particular those associated with training) to help crews cope with the diagnosis and mitigation of increasingly complex failure modes.

The first case study looks at the installation and operation of urine reprocessing components in the International Space Station (ISS) Environmental Control and Life Support System (ECLSS). This subsystem is critical for long duration missions with limited opportunities for re-supply. The ECLSS is intended to produce a purified distillate from condensate, crewmember urine, and urinal flush water. The output is combined with other reprocessing systems in the water processor assembly to support the oxygen generation assembly and to provide crew drinking water. The US Orbital Segment's Water Recovery System was dispatched with STS-126 (Harwood, 2008). It was important to activate the urine reprocessing assembly as soon as possible so that the orbiter could return a sample to earth for analysis. NASA managers hoped to collect bacteriological and taste study data on the urine recycling system for 90 days. This information would then be used to support a dress-rehearsal using the crew of the following orbiter mission to simulate the load imposed on the ISS life support systems when the ISS permanent crew size was increased. A key aim was to mitigate the risks associated with any potential source of illness that could debilitate more than one member of the ISS crew. The water recovery system not only played a strategic role in supporting additional crewmembers for the ISS. It was also intended both to reduce the costs and hazards associated with the resupply missions that would otherwise provide additional water.

The first indications that the Urine Processing Assembly might be operating in a degraded mode occurred when the unit alarmed on the 20th November 2010. This incident raised particular concern because it was associated with the possible release of combustible materials into the ISS. However, the crew followed standard operating procedures and depowered the unit to put it in a safe state i.e. removal of any ignition sources in the vicinity of the potential release. Data was sent back to the ground teams for analysis. Initial concern focused attention on the cooling to the equipment racks. However, it quickly became apparent that this was a false alarm and there was no gas release.

This incident illustrates a considerable degree of resilience and flexibility, not simply in responding to the problems with the ECLSS alarm but also in coordinating the crew's response to multiple warnings. The UPA alarm was triggered at the same time as the crew was also responding to a warning associated with an Extra-Vehicular Activity (EVA). This second alarm indicated a build-up of carbon dioxide in one of the astronaut's spacesuits. The crew member conducting the EVA was extremely fit. He, therefore, metabolized carbon at a higher rate than expected. This began to reduce the capacity of his C02 absorbent canister. Similar alarms stemming from an individual's metabolism had been seen during training. The ground teams suspected that this might be the causes; however, it had never occurred during a mission. There was still a possibility that the alarm indicated a fault with the suit. In consequence, the flight documentation and SOPs placed extremely conservative limits on

permissible CO2 levels**.** Ground control terminated the EVA and sent the astronaut back to the airlock where the crew member could reattach his suit to an umbilical.

This incident illustrates a range of resilient behaviors that the crew used to cope with an uncertain and changing environment. They had to respond to multiple, simultaneous alarms from the UPA distillation assembly and the CO2 monitors. They also had to address communications problems. On the way back to the airlock, the astronaut found it difficult to hear his crewmates and flight control. It later emerged that his headset volume control knob had been inadvertently turned down. This illustrates some of the problems that arise when applying conventional forms of hazard analysis and risk assessment to human space flight. It would be easy to assign an extremely low probability 'degraded modes' to these simultaneous alarms and degraded modes. As we shall see, however, multiple alarms and component failures characterize many different, complex space operations.

Investments in Training: This paper argues that financial pressures must not erode the positive behaviors that enable flight crews and ground teams to successfully i.e. quickly and robustly resolve degraded modes of operation. In particular, it is essential that training budgets are sustained so that groups of co-workers can rehearse the communications and problem solving practices that are a prerequisite for safe and successful operations. In the ECLSS case study, many different people worked together to address the UPA and IVA warnings. The cooperation between the crew and ground teams was built upon a clear allocation of tasks developed and reinforced by pre-flight training. This illustrates the anticipation that is advocated by the proponents of resilience engineering. It should also be noted that EVAs are recognized to be one of the most hazardous operations associated with the ISS. In consequence, significant resources were dedicated to support the crew. The CO2 warning was continually monitored by the flight surgeon and by the EVA console in the flight control room. During the active periods of EVA preparation and execution they were assisted by the EVA Safety Console team in the Mission Evaluation Room (MER) – similar to the manner in which IVA operations also has a dedicated console. The MER continually assessed the risks to EVA crewmembers. The division of tasks helped to ensure that ground support would not become distracted by the UPA troubleshooting that occupied other members of the ISS crew during the CO2 warning.

Investments in SOPs: Funding is also necessary to support mission planning. For example, the immediate responses to the UPA and to the EVA warnings were guided by standard operating procedures (SOPs). These specified a range of tasks that were intended to mitigate potential hazards both to the astronaut outside the ISS and to the crew inside. These SOPs were informed through conventional hazard analysis and risk assessment. They had also been validated through pre-flight testing. For instance, the analysis of previous training exercises had made mission support aware of the possibility of elevated metabolic rates in some astronauts. The response to the initial UPA warning illustrated the way in which SOPs combined with a successful allocation of ground resources to address multiple alarms. The behavior of the crew and mission support illustrated their resilience to simultaneous problems that might otherwise have posed safety concerns. There are, however, other situations in human space flight when it is far more difficult to determine appropriate responses to multiple warnings. This can be illustrated by the tensions that exist between flexible responses to emerging problems and the use of ad hoc, 'work arounds' for degraded modes of operation.

In the aftermath of the initial warning, the UPA continued to operate for almost three hours before shutting down with a further alarm. The Flight Director, therefore, took the decision to suppress the UPA warning. Ground crews were confident that these were spurious warnings. There was also a concern that future alarms might erode the crews' finite perceptual resources. The decision to remove the alarms was, therefore, justified in terms of human factors concerns. This response was also supported by a ground-based risk assessment involving the various teams mentioned in previous paragraphs. However, there was still the possibility that important information might be lost. By suppressing the warnings, new hazards might have been introduced by removing important information from the flight crew. This was a significant concern given that the cause of the UPA failure remained undiagnosed. In other words, the decision to suppress UPA warnings was supported by a human factors analysis and a multi-disciplinary risk assessment. This situation illustrates the complex engineering judgments and programmatic risk trades that must be made during many human space missions; to support the crew by suppressing a potentially spurious alarm or to retain the warning even though it eroded finite perceptual resources with the small likelihood that it might be conveying meaningful information either now or as a result of future failures. The skills and expertise required by both the ground teams and flight crews in making such judgments can only be developed through the careful planning and subsequent exercises that take place in the weeks and months prior to a mission and repeated practice on a regular bases during on-orbit, On Board Training.

Investments in Redundancy and Defenses in Depth: One of the most significant investments in mission assurance comes from the multiple teams that are coordinate the response to major incidents. Each attempt to address the UPA failure was subjected to a detailed analysis by the Safety Console team within the Mission Evaluation Room (MER). They had the power to approve or to block all troubleshooting and maintenance procedures. In order to reach such a decision, they consulted the engineering groups involved in the design and certification of the distillation unit. These different teams worked together to assess the potential hazards associated with proposed interventions. In many other contexts, including the military and air traffic management, operators are left to improvise solutions to degraded modes of operation without additional support from development and maintenance staff. A range of safety monitoring functions protects ISS operations. For instance, the MER Safety Console provides N-2 or 'two-fault' tolerance. In other words, they must ensure that safety is maintained even if there are two simultaneous faults in ground-based or orbiting systems. The redundancy implied by the N-2 approach increases confidence in the systems infrastructure but can only be sustained at significant cost. In addition, the Safety Console continuously monitors ISS operations. They coordinate the hazard analysis that guides subsequent interventions to diagnose or mitigate degraded modes of operation; including the UPA failure. ISS monitoring functions provide significant protection beyond that in most other industries. They offer additional assurance that flexible interventions will not undermine the safety of complex systems.

Investments in a Resilient Mission Culture: Previous sections have argued that resilience in space operations is based on pre-planning. This, in turn, requires significant financial resources in order to create the team structures and SOPs that focus interventions during degraded modes of operation. These pre-mission activities also sustain the working friendships and informal communications patterns that reinforce more formal patterns of behavior. The increasing complexity of many recent space missions arguably reinforces the importance of these investments. Integration has led to the evolution of systems of systems that are supported by multiple levels of redundancy. This provides strong benefits in terms of dependability. However, it also makes it more difficult to diagnose the underlying causes of alarms. Management and engineering staff must decide whether particular warnings pose significant hazards to future operations. If there appear to be no adverse outcomes after a particular warning then there is a temptation to scale back the resources that are allocated to fault finding. This creates significant concerns when undiagnosed problems remain even though they do not appear to undermine the safety of complex operations. In contrast, the mission culture of the ISS teams encouraged sustained and persistent efforts to diagnose the causes of the ECLSS alarm.

The UPA assembly relied on a centrifuge to compensate for the lack of gravity during the separation of liquids and gasses. Monitoring results showed that excessive vibrations were causing protection software to shut the unit down. A number of possible explanations were developed. The first suggested that thermal expansion occurred after the unit had been running for some time. This could account for the friction or blockage that led to the motor symptoms; including a speed reduction and increased current. A second explanation proposed that the distillation assembly reached an operating frequency that caused the unit to move so that a speed sensor came in contact with the spinning centrifuge. A third hypothesis suggested that there were interactions between the previous explanations and other (unspecified) causes.

Both resilience and persistence can be seen in the manner in which ground teams cooperated with the crew to develop different mitigation strategies. One approach was to return the UPA on STS-126 for repair. The distillation assembly was designed to be removed and replaced on-orbit. In consequence, these procedures had already been subject to a safety assessment. So while the operation was unplanned, it was not unexpected nor was it unusual within the context of ISS operations**.** However, the return of the unit would have eroded the time available to test samples before the UPA was needed to support the enlargement of the ISS crew. Further concerns arose because there were no alternate assemblies that could have been brought up once the original centrifuge was removed. A further option was to delay the return of STS-126 in the hope that repairs could be completed on orbit. This would enable sufficient samples to be obtained prior to the Orbiter's return. As these options were discussed, the crew identified a further 'solution'. The aim was to tailor the duration of reprocessing activities so that it did not trigger further warning. The UPA would be operational for short periods of time and then be allowed to cool down. If an alarm were generated after two hours then the process should only be operated for up to 1 hour and 45 minutes before cooling.

Mission control worked hard to minimize safety concerns; hazard assessments were developed for each proposed intervention. Safety and engineering teams continued to work with the crew to identify the cause of the alarms. This led to a further explanation; rubber washers that reduced the noise from the centrifuge might also be allowing sufficient motion to create

harmonic effects. The crew, therefore, tried to minimize vibrations by removing the rubber washers. The centrifuge was 'hard-bolted' onto part of the UPA mounting. The process was then restarted. Ground telemetry showed that the assembly was working normally. However, the crew reported hearing unusual noises from the centrifuge 'as though something was off-balance'. Further contingency plans were developed to extend STS-126 to enable the collection of additional samples, with the possibility of bringing the distillation unit back. The continued operation of the UPA justified the decision to retain the assembly on the ISS.

The crew and ground teams cooperated to develop 'work arounds' for the problems that led to the UPA alarms. The apparent flexibility and lateral thinking demonstrate a host of resilient behaviors. However, the unit failed again after the departure of STS-126. Subsequent analysis identified that the problems were caused by the loads imposed on the distillation unit during launch. This raises the question of whether finite mission resources were wasted trying to identify the cause of the problem. These might have been saved if a less flexible approach had forced the replacement of the distillation unit with the return of STS-126. This argument also relies on hindsight. If the unit had not subsequently failed, we would have applauded the tenacity shown by crew and ground support as they worked to fix a pathological degraded mode during the installation of the UPA.

Previous paragraphs have identified a paradox. Successful space missions rely on careful planning, the development of SOPs and efficient communications practices in order to sustain flexible responses to uncertain events in complex environments. They have also identified the limits of resilience when a flexible response might consume mission resources with ad hoc 'work arounds'. This creates a huge challenge for engineering management as finite resources of time and money are eroded in the iterative refinement of multiple solutions. Undue flexibility has also undermined safety in industries that typically lack the oversight, which protects the ISS (Johnson 2009, Johnson, Kirwan and Licu 2009).

3. THE ISS P6 SOLAR ARRAY DEPLOYMENT CASE STUDIES STS116 AND STS120

Previous sections have described how pre-flight investments in planning and training help to create the team structures that support the resolution of complex and unpredictable challenges in the engineering of human space flight. They nurture the formal and informal communication mechanisms that help to integrate standard operation procedures, for instance based around predetermined responsibilities following multiple alarms, and real-time risk assessments that support more flexible intervention, while trying to repair the UPA. This integration provides both assurance and resilience but it requires a significant budget that must be protected in times of financial stringency, especially when other development costs may exceed initial expectations.

The previous case study illustrated the manner in which pre-flight investments in crew training can encourage a coordinated response to complex failures involving multiple alarms. There are, however, a small number of failures that stretched these multiple defenses for the ISS and its crew. For example, the P6 solar array was damaged on deployment during STS-116. The subsequent tear prevented it from being retracted or extended. This compromised the structural integrity of the array. Similar 'pathological' situations had been considered during mission planning. Even so, this incident stretched the coordination and ingenuity of both the crew and ground teams. The loss of structural integrity created a host of concerns that prevented the Orbiter from undocking.

The partial deployment initially blocked the operation of the Solar Alpha Rotary Joint (SARJ). This prevented the solar arrays on the P3/P4 truss from rotating to follow the sun and created further concerns for ISS power management. On flight day 5 of STS-116 more than 40 commands were issued to furl and unfurl the jammed array in order to remove a number of kinks caused by an apparent loss of tension in the guide wires. After some seven hours of coordinated efforts between the crew and the ground team deployment provided sufficient room for the operation of the SARJ. Further efforts were abandoned as the crews needed to rest. This also provided an opportunity for ground teams to reassess their options.

A number of 'work arounds' were identified to help deploy P6. Many of these were improvised using the insights that had been gained during pre-flight planning. For example, the crew had observed oscillations on some of the solar arrays when they were using exercise equipment. They, therefore, tried to use this equipment to induce further movement in the truss. This was unsuccessful and ground teams continued to analyze the design of the assembly using the problem solving skills that had been employed in the exercises mentioned above. They eventually concluded that an EVA would be required to address the problems. This illustrates further complexities in implementing a resilient approach to degraded modes of operation. As we have seen, EVAs are known to be one of the highest risk operations conducted by the ISS crew. However, the risks during an EVA had to be balanced against the continuing hazards associated with the threat posed by

the structural problems arising from the P6 deployment. It is impossible to guarantee that high-risk interventions will achieve their intended outcome. In this incident, the subsequent EVA only succeeded in retracting a further six bays of the assembly. A further EVA had to be scheduled in order to complete the task. This took just under seven hours towards the end of the STS-116 mission. The duration of the EVA provides an indication of the complexity and also of the risk-exposure associated with this P6 deployment activity.

This incident again illustrates the funding nexus between resilience and pre-planning. The successful resolution of the P6 problem depended on close coordination between the flight-crew and ground support. The operations performed during the subsequent EVAs had not been rehearsed. However, they depended upon skills and expertise that had been developed prior to the mission. This included the risk assessment and hazard mitigation procedures that were validated in simulators and exercises prior to launch. It remains to be seen whether this level of skill and expertise can be sustained when financial pressures continue to affect pre-mission training.

The problems experienced with the P6 deployment during STS116 are not the only example. For instance, STS-120 included a further assembly mission to move the P6 truss segment from the Z1 Node to its permanent location on the P5 truss. This was necessary to enable the subsequent launch of the European and Japanese laboratories. There were unique features of this mission; in particular assembly tasks were scheduled to continue after the departure of the Orbiter to relocate the Harmony module. This process required both careful pre-planning and, as events unfolded, careful improvisation during seven EVAs and many more robotic maneuvers using the ISS robotic arm. The relocation was complicated by the size of the P6 assembly and also the distance that it had to be moved across the ISS superstructure. This led to a plan in which an initial EVA was used to disconnect the electrical and mechanical systems. The robotic arm was then employed on the next day to move the component to the intended destination before a further EVA re-established the necessary connections; “The techniques employed during the P6 installation operations on STS-120/10A were developed following the review of the loads analysis and several evaluation sessions in the virtual-reality training facility at the Johnson Space Center” (Aziz, 2010). The P6 assembly was relocated without significant problems until attempts were made to deploy the solar arrays. This led to further problems with a small tear between two panels being created when the guide wires became tangled. Many of the concerns that were raised during STS-116 now re-emerged. In particular, the structural integrity of the ISS could not sustain the docking or departure of an Orbiter. Power generation was again limited and urgent plans had to be made for the EVAs that would be needed to resolve the problem.

This incident helps to illustrate the complexity involved in planning an EVA in response to such incidents. The robotics flight control team quickly realized that the ISS arm would not be long enough to place a crew member at the site of the damage to P6. They, therefore, began to develop workarounds that would extend the reach of the robotic arm. One approach involved the use of the Orbiter Boom Sensor System (OBSS). This was used to inspect the thermal protection and extended the reach of the arm by around 50 feet. The crewmember conducting the EVA could then be placed on the end of the OBSS providing an additional foot restraint could be placed on the inspection boom. As mentioned before, EVAs perceived to pose the greatest risks to the crew. In consequence, senior ISS program management directed a detailed consideration of the engineering and safety issues. The initial plan to conduct the EVA was postponed for 24 hours in order to enable the development and testing of repair techniques. However, before the plan could be put into effect the robotics team had to coordinate with the other flight control disciplines. The EVA did not simply require the installation of the OBSS; it also included the repositioning of the ISS to improve lighting for the working area. It also included the return of the OBSS to the Orbiter after the repair had been completed. The coordination of the improvised plan was focused on a timeline of milestones stretching well before the EVA. This was then used to identify priorities for more detailed analysis, for the allocation of resources including crew time and also for the associated risk assessments.

The OBSS was not designed to be used an extension for the ISS robotic arm. Ground teams had no experience of the dynamic, kinematic properties that might emerge during the operation of the combined systems. These uncertainties were compounded by the lack of either flight crew or ground team training in the deployment of this improvised system. There was also a pressing need to configure the robotic software so that it could be used to control the movements of the crew member on the end of the OBSS during the EVA. Many of the procedures that were usually employed to validate the frames used to direct the control software had to be abbreviated. However, a plan emerged to use a second crew member during the EVA to monitor the progress of the operation and provide immediate feedback to the rest of the crew. Joint conferences were held between the crew and the flight control team to brief each other on the hybrid operation of the robot arm and OBSS assembly. A review of the successful completion of this

repair identified a number of lessons for future missions (Aziz, 2010). These included the important of fault-tolerant planning and of detailed scripting for procedures that affect interdependent systems. Both of these issues have been mentioned in previous sections of this paper. The closing sections of the review advocated that greater resources be devoted to pre-mission contingency planning; "The damage sustained during the deployment of the 4B solar array caught the ISS program and the flight control team by surprise. Despite problems observed during the retraction of the 4B array on the 12A.1 mission, no one was prepared for the possibility of problems during the deployment operations. As a result, no assessments were performed pre-flight to determine the feasibility and the techniques for positioning an EV crew member at the solar arrays to perform repairs. Performing such assessments prior to the mission would have significantly reduced the valuable time and effort spent during the mission and would have allowed the flight control to develop preliminary products to support this contingency. While the likelihood of the problems observed during the mission may have been considered low prior to the flight, the consequences of those problems were known to be severe enough that some contingency planning for solar array repair should have been performed as part of the pre-mission preparations".

4. THE ISS CENTRAL AND TERMINAL PROCESSING CASE STUDY

A further example of importance of pre-planning in maintaining safety and ensuring a resilient response to degraded modes can be provided by the simultaneous failure of all six Russian ISS central and terminal computers during STS-117. The loss of computational support affected the Russian components Environmental Control and Life Support System (ECLSS). Software systems helped to regulate the ISS Elektron Oxygen generator. At the time of the failure there were also substantial oxygen reserves; up to 56 days for 10 astronauts. There was also sufficient CO2 scrubbing capacity and temperature control for both the U.S. and Russian segments. As in the previous case studies, standard operating procedures and safety management principles again helped to guide the response as the N-2 or dual fault principle was invoked. In this case the ground teams worked to provide an alternative back-up for the Elektron system. A plan was quickly developed and validated to install a hydrogen vent valve during an additional EVA. This enabled a new U.S oxygen generator to be brought on-line.

In addition to the loss of the ECLSS Elektron subsystem, the computational failures also affected attitude control for the ISS. Control moment gyros (CMGs) could be spun to counteract induced momentum during normal operations. However, these were insufficient to control the forces created by disturbances such as an orbiter undocking. In such cases, the CMGs must be taken offline and the station allowed to enter free drift. Once the Orbiter has undocked, Russian computer-controlled thrusters can be fired until control is returned to the CMGs. Without the Russian software for the on-board thrusters, the ISS relied on attitude control from the Orbiter's thrusters. This created a catch-22 situation where the ISS relied on the Orbiter to counteract any momentum imparted when that Orbiter undocked. If the Orbiter could undock then it was likely that the gyroscopes would quickly have become saturated and the only apparent way to avoid a loss of control would have been to fire the ISS thrusters which, in turn, depended on the failed computer systems.

The crew and the ground teams faced a novel set of problems. Some elements had been rehearsed in training others had been addressed in previous missions; for instance relying on the Orbiter for attitude control. Other elements, including the knock-on failures associated with the loss of the central and terminal computers had not been considered before. Additional complexity arose because the problems first emerged during an EVA to repair a torn thermal protection blanket on the port orbital maneuvering system pod of the Orbiter.

The response to the computer systems failure shows strong similarities to the previous UPA case study. However, the consequences were potentially more serious. The Orbiter has been scheduled to return one week after the initial systems failure. The STS-117 crew, therefore, worked to extend the duration of their mission. This included procedures to reduce the Orbiter's power consumption. At the same time, ground teams began to find 'work arounds' for the initial failure. One course of action focused on using thrust from a Soyuz or Progress cargo ships after the departure of STS-117. The development and safety assessment of these plans was again guided by joint procedures developed and rehearsed between Russian and US engineers. However, this incident arguably revealed the need for increasing cooperation in joint exercises to resolve complex infrastructure failures

At the same time as the crew and ground teams worked on restoring attitude control, attention was also focused on the potential causes of the failure. As mentioned above, the computers went down at the same time as an EVA was taking place. One task during this procedure had been to connect a power supply between the Starboard 3 and 4 truss assemblies. The intention was to route power between S3 and S4 to the S6 truss when it arrived. However, this connection was not needed at the point when the computers crashed. Initial

hypotheses considered possible interactions between the station's solar arrays and the service module housing. These interactions included electromagnetic interference or power supply problems. The ground teams, therefore, decided to schedule an EVA that would disconnect the newly installed but unused power supply. Further monitoring of the power systems failed to identify potential causes for the failure. There was an increasing realization that the simultaneous occurrence of the EVA and the computational failure might have been little more than coincidence. Other hypotheses suggested that the increased size of the ISS might be causing electromagnetic charging from the Earth's magnetic field. As with concerns over the power supply connections between S3, S4 and S6, there was a pressing need for scientific and engineering data to support speculation. This provided further lessons for the planning and rehearsal of ESA's Mars500 project. A number of pathological failure scenarios have been deliberately inserted to test flight crews and ground teams, including major power system failures with a twenty minute communication delay and multiple potential causes.

In the hours after the initial failure, the crew worked to restore the systems that had failed. Together with Russian and US ground teams, they were able to test a single channel on two of the failed processors. They were also able to reconfigure the power management systems. However, they were unable to reboot the attitude control systems. These initiatives had to be synchronized with the crews' scheduled sleep periods. The computer repairs also had to be suspended when the ISS moved out of range of the Russian ground controllers. This reinforced lessons for the coordination of future pre-flight training. As mentioned above, the Mars500 project is deliberately replicating the temporal characteristics of communications between the ground teams and the crew as they combat a range of technical failures during the simulated mission.

At this stage, failure hypotheses focused on the power quality issues mentioned earlier. Concern also focused on software failure modes associated with the order in which the two primary computers were restored. This led Russian ground teams to identify problems with the secondary power supplies that supported three redundant communications channels between the failed processors. The crew, therefore, used a jumper cable to bypass one of the channels. This left the remaining two channels functioning correctly. They were then able to boot four out of the six navigation and command systems. Plans were made to bring forward the date of a Progress mission in order to replace the damaged secondary power supplies. Although the computer systems seemed to be working normally by the time that the Orbiter undocked, there was still no clear causal explanation for the failure. Nor was there an agreed explanation for the success of the jumper cables. This led to significant uncertainty. The jumper cables were viewed as a short term fix and engineers were uncertain whether a similar failure might occur in the future even after the secondary power supplies had been restored. In consequence, monitoring tools were used to identify potential problems at different layers in the communications protocols. The crew systematically checked the hardware and network components. This reinforces observations made in previous sections about the long running nature of the UPA problems. In the past, many exercises and pre-flight drills focused on problems that could be resolved over a few hours or days. In both of the examples, engineers had to work with the flight crews over prolonged periods of time during which there was no single causal explanation for the problems that they had experienced. Such extended uncertainty can be difficult to recreate during pre-flight training under significant financial constraints.

The detailed inspection of the data cabling systems helped to identify that there was corrosion on one of the connectors for the BOK-3, secondary power monitoring system. This had been by-passed by the jumper cables. Water condensation was, in turn, identified as the cause of the corrosion. The condensation had been created by repeated emissions from air separation lines that were part of a nearby dehumidifier. Under normal operating conditions, the cabling should remain warm enough to prevent condensation from forming. However, the dehumidifier was itself operating in a degraded mode. It continued to turn itself on and off, thereby generating surges of cold air that reduced the temperature of the computer cables to a point where there was condensation. A design review subsequently identified that the corrosion could trigger a disconnect command across the three redundant channels of the computers power monitoring system. This was intended to protect against unintended power fluctuations. However, it also triggered a common cause failure for the triple modular redundancy used to protect data communications. This illustrates a relatively common situation in which redundancy and extra layers of protection can inadvertently bring down safety-related systems (Johnson, 2009a).

The response to the secondary power supply failure also illustrates the complex forms of risk assessment that must be considered by ground teams supporting crew interventions. In this case, the jump leads reduced the risk of a common mode failure across the multiple redundant communications channels. At the same time, however, it exposed computational systems to any power surges that would not have been screened by the monitoring systems. By rerouting the power monitoring systems, the Russian computer systems continued to

control the ISS thrusters until STS-118. The Orbiter was then used to provide attitude control while the crew replaced the faulty power units. The replacement procedure provides further illustrations of the work arounds that characterize the engineering of complex systems. The members of the crew discovered that one of the cables was 40cm too short to replace the existing section of the power monitoring system network. In consequence, the original cabling had to be retained after a further visual inspection had determined that it was not corroded. The MER Safety Console helped to coordinate approval for the original cabling to be retained before the jumper cables were removed.

5. CONCLUSIONS

The ISS Program is designed to deliver mission success, including crew safety. It does this through an approvals process whereby individuals assume responsibility for specific decisions. For example, the ISS Program Manager must sign to acknowledge they have understood the consequences of a safety non-conformance report. These approved processes are implemented at "arms-length" from the ISS Program management. However, financial constraints may result in two organizational responses:

- A curtailing of these approval process (or in the worst case pressure to circumnavigate existing checks and balances through particular "work-arounds");

- A reduction in the resource / expertise levels that execute the safety approvals processes.

For safety and training (by way of example) resilience engineering approaches should assure that there is a robust application of the approvals processes and that expert resource is made available to engage in and inform these processes.

This paper addresses the engineering and operational consequences of growing fiscal pressures on human space flight programs. In particular, we have argued that the budgets associated with pre-flight planning should be protected as much as possible. Our arguments have been based on an analysis of previous interactions between flight crews and ground teams in response to systems failures and degraded modes of operation. It is clear that pre-flight planning makes three principle contributions. Firstly, it helps prepare the organizational structures that are necessary to respond in time critical situations. As we have seen multiple simultaneous failures stretch finite crew resources and, typically, require the level of coordination that cannot be achieved without significant practice. Secondly, pre-flight planning provides opportunities for the validation and refinement of standard operating procedures. Again, these are critical because there may not be time for a detailed risk assessment when interventions may be necessary in minutes rather than hours. SOPs provide a framework for intervention that can be tested in drills and exercises, providing an opportunity for failure before the crews' lives are at stake. Finally, pre-mission planning increases resilience by establishing the informal relationships that support effective communication under degraded modes of operation. This is particularly important given that it is impossible to predict and train for every possible contingency before a mission starts.

Pre-mission planning helps to guide the allocation of tasks and responsibilities in response to urgent operational requirements. This is apparent in the way in which the ISS crew were able to cope with multiple simultaneous alarms during the UPA warning. Some of the communications problems that emerged between US and Russian teams during the computational systems failure also, arguably, illustrate the need for greater pre-mission cooperation. Exercises and drills based on previous operational scenarios help to ensure that necessary tasks are not omitted or unnecessarily duplicated. This cannot easily be coordinated in the immediate aftermath of an adverse event without repeated rehearsal. In practice, the allocation of tasks has been guided by experience that stretches back to Apollo and beyond. However with the establishment of the ISS, the introduction of multi-national crews and engineering infrastructures creates new tensions and opportunities. As we look to the future, it is also likely that significant changes will have to be made in the division of responsibilities between the flight crews and the ground teams. Many of the long duration mission scenarios will incur significant communications delays. In these situations, there may not be time for the crews to refer critical interventions down to the EVA Safety Console team in the Mission Evaluation Room (MER). The Mars500 exercises are just beginning to provide evidence of the range of problems that could be identified during pre-flight planning for the next generation of human space missions.

The case studies have also shown that pre-mission planning helps to form the Standard Operation Procedures (SOPs) that protect safety. This is important because exercises and drills are not always successful and failure provides the feedback that is necessary to refine working practices before operations occur in the remote ISS instance. SOPs are essential to provide a timely response to adverse events because there is often insufficient time to improvise interventions with the limited resources and multiple hazards associated with human space flight. In extreme contingencies, it must be possible to ensure a "safe haven", i.e. the ability in any instance to isolate the crew from the hazard and its

effects. Some cases this is donning Personal Protective Equipment (PPE) and others retreating to the escape vehicle (Soyuz).

The validation of procedures and processes is critical, prior to launch because multi-national crews must coordinate the actions in ways that can be particularly difficult to negotiate in response to a system failure. This is especially important where the consequences of any intervention can have knock-on effects for different infrastructures provided by different nations. It is clear from our case studies that many of the incidents we studied have pushed the adequacy of SOPs to their limit. For instance, crews often cannot find a relevant procedure for multiple system failures that were not anticipated in pre-flight planning. However, these predetermined procedures provide a common point of reference that guides more flexible interventions that are often necessary in human space flight after an immediate 'safe state' has been achieved.

The first two benefits from pre-flight planning help to create the static organizational and procedural structures that are essential for time-limited responses to adverse events. In contrast, the final benefit of pre-mission planning is that it promotes the flexibility and resilience required when these static structures are insufficient to address systems failures. By repeatedly refining their response to different scenarios in drills and exercises, crews and ground support learn to cope with uncertainty. They develop cooperative problem solving strategies and they learn to make risk-based decisions in areas that are not covered by SOPs. In particular, senior management develops the courage to try a solution and be prepared for it to fail, providing the crew are not exposed to undue risk. This is essential if missions are to continue when the causes of a failure are unknown. This is illustrated by the decision to undock the Orbiter from the ISS even though there was no clear understanding of why the jump leads had enabled the crew to reboot the Russian computational systems. There was still a possibility that an undiagnosed fault could have returned to compromise the attitude of the space station before the next Orbiter mission. However, this risk was identified and acknowledged. It forms a contrast to some of the hazards that were arguably inadequately addressed prior to several of the accidents that continue to haunt human space flight programs and, which may continue to haunt us if we do not preserve the budgets necessary for effective pre-mission planning. Learning to cope with engineering uncertainty is likely to remain a key issue in future operations. The integration of complex systems developed by multiple nations and by different commercial contractors will have significant consequences for the diagnosis of infrastructure failures in future flights. In particular, it is critical to provide open and easy access to cross-platform engineering documentation. It is also important to identify the opportunities that arise from the exchange of information and techniques between International Partners. For instance, when considering introducing resilient training the potential for collaborations between partners ranges from establishing common training and certification standards, to using novel applications of (existing) terrestrial techniques for example "serious gaming" technologies as part of ISS partner training programs.

Acknowledgements

The work described in the paper has been supported by the UK Engineering and Physical Sciences Research Council grant EP/I004289/1.

References

S. Aziz, Lessons learned from the STS-120/ISS 10A robotics operations, Acta Astronautica, (66)1-2:157-165, January-February 2010.

European Space Agency Knowledge Engineering Office, September 2009. Last accessed February 2010, available on: http://www.esa.int/SPECIALS/Space_Engineering/SEMTH8KIWZF_0.html

W. Harwood, STS-126/ULF2 Mission Archive, CBS News/Kennedy Space, November 2008. Last accessed February 2010, available on: http://www.cbsnews.com/network/news/space/126/STS-126_Archive.html

E. Hollnagel, D.D. Woods and N. Leveson (eds.), Resilience Engineering: Concepts and Precepts, Ashgate Publishing, London, UK. 2006.

E. Hollnagel, The ETTO Principle: Why Things That Go Right Sometimes Go Wrong, Ashgate, Farnham, UK, 2009.

C.W. Johnson, Degraded Modes and the 'Culture of Coping' in Military Operations: An Analysis of a Fatal Incident On-Board HMS Tireless on 20/21 March 2007. In J.M. Livingston, R. Barnes, D. Swallom and W. Pottraz (eds.) Proceedings of the US Joint Weapons Systems Safety Conference 2009, Huntsville, Alabama, 3511-3521, 2009.

C.W. Johnson, The Dangers of Interaction with Modular and Self-Healing Avionics Applications: Redundancy Considered Harmful, In J.M. Livingston, R. Barnes, D. Swallom and W. Pottraz (eds.) Proceedings of the 27th

International Conference on Systems Safety, Huntsville, Alabama, USA 2009, International Systems Safety Society, Unionville, VA, USA, 3044-3054, 2009a.

C.W. Johnson, B. Kirwan and A. Licu, The Interaction Between Safety Culture and Degraded Modes: A Survey of National Infrastructures for Air Traffic Management, Journal of Risk Management, (11)3:241-284, ISSN 1460-3799, 2009.

C.W. Johnson, L.L. Fletcher, C.M. Holloway and C. Shea, Configuration Management as a Common Factor in Space Related Mishaps. In J.M. Livingston, R. Barnes, D. Swallom and W. Pottraz (eds.) Proceedings of the 27th International Conference on Systems Safety, Huntsville, Alabama, USA 2009, International Systems Safety Society, Unionville, VA, USA, 3047-3057, 2009.

C.W. Johnson, A. Herd and M. Wolff, The Application of Resilience Engineering to Human Space Flight. In H. Lacoste-Francis (eds.), Proceedings of the Fourth International Association for the Advancement of Space Safety, Huntsville Alabama, NASA/ESA, Available from ESA Communications, ESTEC, Noordwijk, The Netherlands, ISBN 978-92-9221-244-5, SP-680, 2010.

J. Leonhardt, E. Hollnagel, L. Macchi, B. Kirwan, A White Paper on Resilience Engineering for ATM, EUROCONTROL, Brussels, Belgium, 2009. Available on:
http://www.eurocontrol.int/esp/gallery/content/public/library/A%20White%20Paper%20Resilience%20Engineering/A_White_Paper_Resilience_Engineering.pdf

NASA, Procedural Requirements for Mishap and Close Call Reporting, Investigating, and Record keeping. NASA Headquarters, Washington DD, USA (NPR 8621.1B), 2006.

D. Woods, Creating Foresight: How Resilience Engineering Can Transform NASA's Approach to Risky Decision Making David Woods, Testimony on The Future of NASA for Committee on Commerce, Science and Transportation, John McCain, Chair October 29, 2003.

D. Woods, Creating Foresight: Lessons for Enhancing Resilience from Columbia. B. Starbuck and M. Farjoun (eds.), Organisations at the limit: Learning from the Columbia Accident, Blackwell, Oxford, 2005.
http://www.esa.int/SPECIALS/Space_Engineering/SEMTH8KIWZF_0.html

W. Harwood, STS-126/ULF2 Mission Archive, CBS News/Kennedy Space, November 2008. Last accessed February 2010, available on:
http://www.cbsnews.com/network/news/space/126/STS-126_Archive.html

E. Hollnagel, D.D. Woods and N. Leveson (eds.), Resilience Engineering: Concepts and Precepts, Ashgate Publishing, London, UK. 2006.

E. Hollnagel, The ETTO Principle: Why Things That Go Right Sometimes Go Wrong, Ashgate, Farnham, UK, 2009.

C.W. Johnson, A Handbook of Accident and Incident Reporting, Glasgow University Press, 2003. Available from http://www.dcs.gla.ac.uk/~johnson/book

J. Leonhardt, E. Hollnagel, L. Macchi, B. Kirwan, A White Paper on Resilience Engineering for ATM, EUROCONTROL, Brussels, Belgium, 2009.

NASA, Procedural Requirements for Mishap and Close Call Reporting, Investigating, and Record keeping. NASA Headquarters, Washington DD, USA (NPR 8621.1B), 2006.

D. Woods, Creating Foresight: How Resilience Engineering Can Transform NASA's Approach to Risky Decision Making David Woods, Testimony on The Future of NASA For Committee on Commerce, Science and Transportation, John McCain, Chair October 29, 2003.

D. Woods, Creating Foresight: Lessons for Enhancing Resilience from Columbia. B. Starbuck and M. Farjoun (eds.), Organisations at the limit: Learning from the Columbia Accident, Blackwell, Oxford, 2005.