1
A Story of a Failure

"Listen to the people doing the work...That's how managers learn..."

Management's Summary: A worker drilled into the top of a 5 gallon Liquid Propane tank with a ¾ inch drill causing a potentially fatal near miss of an explosion in a pressurized gas storage area.

Management's Action: Discipline the employee in question; send all shop personnel to additional training; rewrite shop procedures to include a section prohibiting drilling into pressurized 5 gallon Propane tanks with hand tools.

Freon comes in many different sized pressurized containers for many different reasons. Once used, any sized empty tank becomes a problem to remove from the facility (even before used, Freon brings with it its own set of environmental issues). The problem is that while the tanks are still under pressure, these tanks cannot easily, legally, or environmentally be removed. In fact, what is normally done is that the cans are punctured, flattened, and recycled.

For this story, the Freon tanks in question are 5 gallon containers. These Freon tanks are used all over the facility and, in a stroke of organizational and environmental brilliance, a central collection point was created in order to provide a place to gather all the empty Freon tanks when they had been used. This central empty Freon tank dumping spot was basically a square pin made of chain-link fencing with a gate. This area was specifically designated as a place for the facility to collect the empty Freon tanks, depressurize and flatten these tanks, and prepare them to be recycled.

This Freon tank collection system worked fairly well. No one person at the company owned the Freon process. The collection area and the process of recycling the Freon tanks was the "brainchild" of a team of workers who had been tasked with creating a way to recycle these tanks, minimize waste, and reduce the costs associated with disposing of the tanks. The maintenance organization kind of inherited this tank disposal program. It was not assigned to the pressurized gas department, nor was it directly assigned to the waste minimization teams. This program was given to the folks who managed the maintenance shops.

When the cage filled up with many empty tanks, someone was selected to go out to the storage area, sort through the tanks, depressurize and flatten the tanks, and pack them up on a pallet to go to the metal recycler. No one worker owned this job. This job was seen as a straightforward process and was usually assigned to new or less experienced workers. It was a straightforward task, requiring little experience, easy to do, and was seen as more of a nuisance than actual work.

For this event, a new worker who was not getting along well with coworkers within the shop was selected to clean out the Freon cage. He was taken to the work area and shown how to pop a small plastic seal on each Freon tank to be recycled, in order to depressurize and begin the process of flattening. He was shown the work activity by a supervisor, observed doing the work twice on his own, and left to finish the project alone in the work area. The worker seemed to understand the task, be able to accomplish the task with few problems, and capable of continuing with the task to completion.

The disposal of Freon tanks is an easy, unimportant, low risk task. This task was considered so unimportant that there was no process owner, no formal procedures, no identified risk assessment activity, and, therefore, no risk reduction activities created to do this work. Why would a worker need any of these things? Recycling Freon tanks is not a high risk activity. Well, that was the case at least until the process failed.

And fail it did. The worker drilled a hole in the side of a 5 gallon, pressurized Propane tank instead of a Freon tank. Propane is explosive. Freon is not explosive. Luckily, in this case the tank did not explode. In this case, the organization got lucky. This is a case of an event that did not happen, and in turn was a gift to the organization since they could learn from it without having a terrible failure.

However, the decision was made by the post-event investigation that this worker should have known the difference between the two types of tank. If the worker had recognized the difference between these two tanks, the worker would not have tried drilling a hole in the tank of explosive gas, placing his life at risk, and, as importantly, the future of the facility at risk. Nothing stops work at a company faster than blowing up a worker.

THE SECOND STORY OF THE SAME FAILURE

At this company's facility there was no disposal path for small, 5 gallon Propane tanks. You know these tanks pretty well. Every Propane grill in almost every back yard in the country has a small, white, 5 gallon Propane tank underneath. Many work groups at this facility owned Propane cookers. These grills were used for cookouts and picnics that were held in the same facility: worker

celebrations, parties, and picnics. The Propane tanks on the cookers last quite a while.

Nothing is exceedingly unusual about Propane tanks and grills being used at a facility. The problem at this company started when the tanks became empty. In this company, it was immensely complicated to replace a Propane tank, or refill a Propane tank. A task that would take minutes to do if you were at your home was almost bureaucratically impossible to do at this facility. This created the problem of having empty Propane tanks at the facility, and no real way to refill these tanks. In fact, at this facility it was much easier to buy a new tank than to refill an old tank. That left empty Propane tanks with no disposal path available.

Unfortunately, the facility had no way to dispose of all of these tanks. There was no Propane tank recycling program. The pressure shop did not handle Propane tanks at any time for any reason. If you had an empty Propane tank in your work area, it was a problem. In fact, the pressure shop foreman said in an interview that an empty Propane tank at this facility was a problem for life. There was no process to dispose of these empty tanks.

Which led to an interesting part of this story of failure. It seems that because empty Freon tanks had a special little fenced-in area, an occasional Propane tank would not be seen as out of place there. After all, in the mind of a worker, a place to recycle a pressurized gas tank is a place to recycle a pressurized gas tank. In workers' minds all tanks are tanks when they have to find a place to deposit an empty one. So, after hours, when nobody was watching, employees at this facility were dropping off empty Propane tanks in with the Freon tanks. These tanks don't look alike, but they don't look that different either. The problem is that the tanks are entirely different. Freon tanks can be smashed and recycled. Propane tanks are almost impossible to smash and depressurize. They are actually designed to be refilled, not recycled.

To add to the complexity of this case, pushing in on a pressure-release seal button on the tank easily depressurizes a Freon tank. Propane tanks are not easily depressurized: even if you open the valve at the top of the tank all the way, it will not depressurize.

If your job is to clean all the white-painted tanks out of the storage area and depressurize, flatten, and recycle all of these tanks, you might find yourself in a position where in order to accomplish the boss's task you must adapt to any and all tanks that appear in the empty Freon storage area. More amazingly, the fact that this worker drilled a hole with a ¾ inch drill through the side of a Propane tank is quite remarkable. It is hard to drill a hole in a pressure vessel like a 5 gallon Propane tank. It takes time and discipline. It is not an easy task to accomplish.

THE STORY OF THE STORY

Suddenly, an event that was thought to have lots to do with a lack of experience, skill, and intellect on the part of this worker (how can that guy be so stupid? Geez nobody drills a hole in a Propane tank) was now a story about the procedure, ability, and policy for the disposal of 5 gallon Propane tanks at this company.

It is easy to say that the worker made an error in judgment in trying to recycle the wrong kind of tank, and, in fact, the worker did make an error in judgment of the situation; but the problem was not the worker's ability to know the difference between a Propane tank and a Freon tank—the problem is much more systemic, and more compelling.

The problem was a unique mixture of some pretty normal conditions at this facility. We can list many of them with little effort: a dual location of many small tanks in the same area, the inability to refill tanks at this facility, the inability to get rid of old and empty tanks, a new employee, the absence of 100% supervision, production pressures, performance management, HR—and the list can go on and on and on.

It is at this point that the facility will feel a need to make a decision. The facility will probably not know they are not making the decision they think they are making. The decision is not about how to handle the worker. The decision is far more about how the facility chooses to see this failure.

THERE ARE TWO CHOICES

Fix the worker (training, discipline, or termination) who did something he did not mean to do, in the hope that he won't again do something he did not want to do in the first place.

Fix the system that determines what should happen with the empty Propane tanks. A system that was clearly not designed to imagine the presence of Propane tanks where Freon tanks were to be stored.

Fixing the worker gives the impression of an immediate solution to the problem, but probably fixes the wrong thing. Punishing the worker is a fast and easy way to "solve" the problem, with the only issue being that it unquestionably fixes nothing at all, not even the worker in question. Because the whole failure will inevitably happen again with a different worker. It may take some time to drift to this point again, but over time it will happen.

Fixing the system that determines where empty tanks are stored and disposed of looks at a much larger problem that will take longer to fix, cost more money most assuredly, and take longer to complete. However, it will fix the right problem, and will ensure that the facility will never be in a position again to accidentally drill a hole through a Propane tank.

The second answer will be better for the worker, better for the organization, and better for the world, and will ensure a safer, more reliable path forward for this organization.

Sadly, it is *not* the path that most facilities normally choose to follow.

The goal for this book is to provide a usable road map for understanding a different way to manage safety in your workplace. In a way, this book should be more like a road map than a textbook. This book should answer one question: "All these ideas are great, but what should I do differently in my organization to make this change? How can I find, and solve these problems, before they become accidents?"

It is essential to note early in this book that I am standing on the shoulders of giants in this field. So much effort and thought have gone ahead of all of us in just understanding the theoretical underpinnings and creating a "new view" for the world of Human Performance and reliability. Many, many times I will direct you through quotes, cartoons, models, and ideas that are directly attributed to these fathers and mothers of the Human Performance idea. We won't spend much time on theory (although it is my hope that this book will make you want to read everything ever written about the topic), that is not the goal.

This book will give you practical guides, tips and tools, and suggestions for managing risks, hazards, and safety.

Practical means that the ideas and techniques you read in this book can be used in your organization almost the day you read them. It is the goal of this book to give simple and fast options you can deploy on the "plant floor" immediately. If you try some of these ideas, you will see changes immediately—seriously, immediately.

This is a substantial claim, a claim that in most cases I would be afraid to make to you. Not here, trust me.

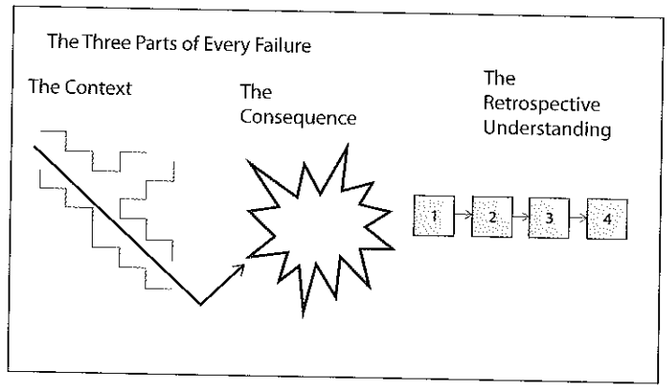

Figure 1.1 The Three Parts of Every Failure

WHAT IS THE DIFFERENCE BETWEEN THE WORD "EVENT" AND THE WORD "FAILURE"?

In our business we talk a lot about failures. Not only do we talk a lot about failures, we talk even more about events. In many ways, these words are used to mean the same thing—they are semi-interchangeable. In many other ways, these two words are worlds apart. We will need to give these two terms an understood meaning for our discussion, but first let's define what we mean by the word "performance," as in the title "Human Performance." Because to understand what is happening, a failure or an event, we must first understand how the idea of *performance expectation* sets the goal. Operational outcomes affect the meanings and context of a word. Performance, by its precise nature, is a word that is defined by its outcome.

WHAT IS PERFORMANCE?

> Performance is the degree to which you get what you expect from a person, a machine, or a process.

Anything other than what you expect is some type of deviation from what you wanted. Sometimes systems and people over-perform (although that seems rare, sadly), and sometimes our systems, and people underperform (that one you probably know intimately). Either way, performance is defined by the degree to which an action or process meets a desired outcome.

In reality, I don't like the terms "event," "accident," or even "failure" all that much. The best way to think about these "distinct moments in time or a process" that have either attracted or should have attracted organizational attention is as a moment where performance expectations were not met. Instead of using "event" or "failure" as words to describe these times, I actually prefer to use the term "deviation from an expected outcome."

Any performance expectation that is not met (over-performance or underperformance) is a deviation from an expected outcome. An accident, an event, even a near miss, is certainly a deviation from an expected outcome. A process improvement, a good idea, a success against all odds is also a deviation from an expected outcome. One is a bad thing, the other is a good thing, but they both are something different from what was expected.

DEVIATION FROM EXPECTED OUTCOME (DFEO)

The only problem with using the term "deviation from expected outcome" is that it gets really old, really fast. I tried to make an acronym out of this phrase

(DFEO); I never could get that to stick. DFEO just never caught on. I am not surprised it did not stick. It is awkward to say, and difficult to remember. It also doesn't sound very practical. Saying "deviation from an expected outcome" is something that a college professor would say, not a safety professional or even a manager.

So what is the difference between the term "event," and the term "failure"? The quick answer is that one of these terms has a big and bad outcome. The other has a small and simple outcome. But honestly the answer is that there is no difference at all in the words themselves. The difference in the meanings of these terms is in the person using the terms, and in the people that are listening to the terms. There is no difference in the terms, themselves—they are both deviations from expected outcomes. Both of these terms are indicators of a difference in expected performance.

A NEW DEFINITION OF SAFETY

For all too long now our organizations have tried to define successful safety programs by counting the number of times our workers screwed up or got injured. In actuality, what we count is the absence of accidents. Our safety incentives are measured and built around a zero accident goal. The Institutional safety numbers we report are usually numbers of injuries or days away from the workplace. These numbers are usually reported "up and out" of the organization, to customers, regulators, and the world.

You don't have to be a genius to know that something seems oddly wrong about the way we measure safety success. We count the number of people we hurt, and totally discount all the people we are keeping safe. The problem is, and always has been, you can't count what doesn't happen. It is hard to count the millions of decisions that are made every day in the field that don't lead to some type of failure. Those millions of operational decisions were all safety and performance reliability successes.

This new definition of safety or performance reliability demands not only a different way to think about events, accidents, and failures, but also, more importantly, a new way to manage your organization. You must change the way you manage safety in order to align your organization and operations to a new definition of what "safe" is, and why it matters. Assuming your management knows how to lead and manage this new method is a gigantic mistake, and will lead to horrible outcomes. These are new ideas and new ways of doing business. How would your managers know this information if you did not inform them of this new way of doing business?

Management now wants the safest workplace that is humanly possible. Obviously, no manager wants workers to get hurt. Knowing that management's intent is true, good, and hopeful, allows us to search for reasons for not changing other than motivation.

Those reasons almost always come back to one single factor. Most managers want to manage safety and performance using the best and latest ideas for success. The problem is that we, as a new-view safety professional community, have not told them about those ideas, and how to use this new view. Take the way we define safety—this is a perfect example of the old world view dominating our new world organizations.

SAFETY DEFINED

> Safety is not the absence of events; safety is the presence of defenses.

Safety can't be the absence of events. Safety is the presence of defenses within the system, environment, and processes. If you constantly choose to use the traditional standards in order to measure safety success, you will constantly create workplaces where information about events and injuries is hidden from management. Worse yet, we are measuring our failures.

The problem is that it is hard to measure what doesn't happen. For every accident that occurs at a workplace, how many accidents didn't happen? How many people didn't get hurt? How many potential problems did workers who identified the error precursors before they had a chance to become full-fledged events solve? These numbers are seemingly impossible to count. That could be why these numbers are not valued by your organization.

In reality, safety is probably best defined by this idea:

> Safety is the ability to perform work in a varying and unpredictable workplace environment.

This definition, although probably the most accurate, is the most difficult to use operationally and discuss with the workforce. Yet, we do this every day. Every day we live in an unpredictable and constantly varying world.

HUMAN ERROR

It is distressing enough that the world is varying and unpredictable. Human beings are also prone to their own unpredictability. Humans make errors. People are fallible, and even the best of us make mistakes.

Error is an unexpected deviation from an expected outcome. Error is an unintentional event. Error is doing something that you didn't intend to do. Errors happen all the time. Everybody makes errors, everybody. The very worst performers make errors. The very best workers make errors. Error is a predictable and natural part of being a human being.

People make a lot of errors. Not all errors have a consequence. In fact, not all errors are actually errors. We only really notice an error if it has some type of outcome or consequence that is large enough to be noticed by either you or other people around you. Error only becomes apparent if you notice an error.

Contemporary wisdom says that the average skilled worker, workers who work with their hands, makes 5 to 7 errors per hour. That same wisdom says that a knowledge worker, workers who work with ideas and concepts, not making things, makes between 15 and 20 errors per hour. True or untrue, we all make errors. Errors are how we are wired, how we are made, a natural part of being human. Human error is inevitable—all workers are error-making machines. What all this means is pretty simple: error is everywhere, and there is nothing you can do to avoid the errors. You can't punish error away. You can't reward error away. Error is an unintentional, unpredictable event. You know it, and I know it. Now we have to make sure that our organizational managers remember that fact.

But it is not that easy. Error is always attributed in retrospect to the worker by the organization after some type of consequence happens to the organization. In other words, the worker making an error is determined after an event happens, not before, and is usually seen as a moment in time where the worker did something wrong. If the worker had done something other than what the worker did, the event would not have happened.

This is always going to be true, but wrong, of every failure that will ever happen in your organization. The notion that the worker chose to make this error is also always true, but wrong. Choosing to make an error only becomes an actual choice after the event has happened. In fact, error can never be a choice. Error can never be a violation. Error is simply the unintentional deviation from an expected behavior.

It is easy to find errors in retrospect. It is even easier to judge these errors as wrong in retrospect. This process is called a fundamental attribution error. Think of it like this:

> A worker is walking across the office parking lot to go in to his office. During this walk, the worker steps on a rock and sprains his ankle. The injury is bad enough that the worker has to go to the clinic and have his leg treated. The worker ends up with a series of x-rays, a cast on his foot, and a day or two off from work.

This is an example of an error; however, this error will be attributed to the worker's judgment and walking ability. At some point, some manager will comment that if this worker had "watched where he was going" this event would not have happened. If this worker had cared more, the worker would have been more attentive. If the worker had been more attentive, the worker

would have stepped over the rock, and completely avoided this injury. This injury is clearly the worker's fault.

The organization is attributing the error to the worker's judgment and behavioral choices. In a way, what the organization will do is assume that a bad outcome must happen to a bad person. You will hear some manager say something like this, "if only...if only the worker would have paid more attention..."

Read on and see how remarkably unfair that way of thinking is to your workforce.

> The next day, the manager of the "twisted ankle worker" who was injured the day before was walking across the same parking lot. The manager had a lot on her mind. She had a reportable injury. She had a worker who had to go to the hospital and get medical attention. She had an employee that had gotten hurt under her watchful eye. While this manager was thinking of all these things she stepped on a rock, the very same rock, in the parking lot, and twisted her ankle. She was in pain and could barely walk...but she didn't report the injury.
>
> Instead, this manager found out whose job it was in her organization to sweep the parking lots and sidewalks. The manager immediately called the roads and grounds crew supervisor, and had his crew sweep the parking lot. Because this manager was so busy thinking about the event that had happened to her worker, she was concentrating on her safety problem and not on where she was walking.

What is intriguing about this concept is the idea that when the worker walked incorrectly it was the worker's fault. When the manager walked badly, her response was to not only fix the problem, but also to move the fault from her as the manager to the person whose job it was to keep rocks off the parking lot. That is the fundamental attribution error, and it happens all the time with safety events. Every critique or after action meeting you will ever attend will have a moment when some manager will claim that had he been the worker he would not have made the same errors that were made by the worker in question.

Human error is a weird thing. Error is always present, but not always noticeable or noteworthy. Error is usually attributed after the fact—a way to explain rationally what happened and why. Error is hard to explain and hard to predict—these are accidental actions done without knowledge of the consequences. Error is never intentional—there is no such thing as an honest mistake because there is no such thing as a dishonest mistake—they are all errors, unintentional deviations from expected behaviors.

If all this is not complicated enough, the real question is if an error is actually ever an error. Error is essential to our discussion, as you will see as you read on in these chapters. Error, however, is not everything—I would caution you not to fall prey to the idea that if you just could stop error you would stop failure. You can never stop error, because errors we identify become visible only after the failure.

IS ZERO ACCIDENTS THE RIGHT GOAL?

The quick answer is you don't want to work in an organization that has any other goal than zero. Zero is probably the best individual goal that any worker could have. What is more compelling is the fact that most workers actually meet that goal. Most workers will work their entire careers and not have a serious accident or event. Zero accidents actually becomes in reality an outcome, the right outcome, but having zero accidents does not make a safety program. It is, at best, the target at which your program is aiming. Asking workers to have zero accidents does not tell workers how to have zero accidents.

The problem is that "zero" is not a particularly realistic goal for the entire organization. In short, the problem is that "zero" is a standard of perfection. Your workers and managers are not perfect, and will make mistakes. They know this fact, and you know it. You will have failures. You will not be able to hold events to a "zero standard" over the lifetime of the organization. You will have failures.

In many ways, that's why your organization has learned to count the number of failures, and track that number—with the goal of constantly driving that number down to zero. It is a noble and noteworthy goal, zero. The only problem with having zero for a goal is that you (and every other safety professional in the world) are only one ankle sprain away from being the worst safety manager at all times.

There is another problem, and that is the problem with using "perfection" as our goal. Perfection leaves no room for the inevitable human error that will happen. And we know it—errors will take place. We just need to learn where and when they could take place. When your organization uses only perfection as its standard, anything other than perfection is just fundamentally wrong. Let's not belabor this point too much; you live this every day of your life, but suffice to say that if a victory in your business is no accident, we will never actually know how you won the contest.

Humans and organizations learn by trial and error. We learn by becoming smarter after we fail. In fact, when we fail, we most often go right back out to the worksite and try again—only this time we try to fail smarter. When we devalue failure, in many ways we are inadvertently devaluing our ability to learn.

WE CAN LEARN A LOT FROM THE AUTO INDUSTRY

Almost every day in some workshop someplace, I ask this question: "How many people were killed in fatal highway accidents last year?" It is a depressing question to ask a group of workers. Every highway death is a tragedy. But this is a compelling question.

Why is this compelling? Partially because in the United States the fatal highway numbers have stayed right around 38,000 people a year for almost 50 years. In 2000, our numbers were around 38,000. In 1990, the numbers were around 38,000. In 1960, the numbers were also nearly 38,000. We have leveled out in the number of people killed on our highways.

Yet, the number of cars on the road has dramatically increased. The number of drivers has increased. Also, the number of millions of miles driven by American drivers increases every year. Still our highway fatality numbers stay essentially the same.

Why?

Every day, workers tell me the answer is that cars have changed. Cars are safer today than they ever have been in automobile history. Today, roads are safer, better designed, better engineered, better marked, better patrolled than in the past. Everything in our highway and vehicle system is better engineered, smarter, and safer. In fact, the auto industry has done everything but change the driver.

The Statement "change everything but the driver..." is important. In many ways on our job sites, we have taken the opposite approach. We have tried to get to safety performance by "leaving everything the same except fixing the worker." It is about time to admit to ourselves that we have been managing safety wrong. We were simply doing safety management backwards.

We can learn much from the idea that our organizations respond in the opposite way. We have almost always "changed nothing but the worker." We rarely fix the system around the worker. We almost always try to "fix" the worker—as our sole corrective action.

Why is it safety programs in our organizations are so interested in fixing the worker and not truly engaged in creating survivable space in our work systems and processes? I fear that the answer may be either we are lazy, or we are cheap. All defenses in vehicle transportation that can be engineered to create the least number of failure consequences possible—or the most survivable space possible—are being created every day with one goal: "We cannot fix drivers, drivers will never be perfect, drivers will screw up. We must design and create systems around the drivers and passengers that keep them safer."

Drivers, according to the daily discussions I have with workers and managers in classes, and workshops, are getting worse. There are many things that make drivers less attentive, less effective, less safe in their operation of a vehicle. Cell phones, texting, GPS units, fancy stereos, iPods—the list can go on and on of things that steal attention away from the highway. In fact, the automotive industry assumes that drivers will be placed in positions of trading attention between driving and "functioning" in the driver's seat.

This is most compelling, because highway fatalities are getting rarer and rarer. In fact, chances are extremely high that you will not die while driving or riding in a vehicle. This is a direct result not of trying to stop car accidents (which are all bad outcomes and should be avoided), but, in fact, of knowing that when an accident inevitably happens the goal is to reduce the consequences of the accident. Make cars fail safer.

Making our cars fail safer? That is exactly the primary idea of defending against accident consequence. Not one person has the power to stop all accidents in the workplace. Human beings simply are not particularly good at predicting the future. I will gladly admit that I am not smart enough to avoid all accidents in my personal life, let alone in my facility.

Because you cannot prevent all accidents, you must assume that accidents will happen, and use your time, energy, effort, and resources in dramatically reducing the consequences of the accidents that will happen in your workplace. You must build systems that allow our workers to fail safer. Start thinking like the car industry, and make your systems safer.

RUMBLE STRIPS

Probably my personal favorite defense against vehicle accident consequence is the humble rumble strip (pardon my rhyme). Rumble strips are the grooves that are placed in the actual road surface along the roadside, just outside the normal driving system, that alert the driver that they have drifted out of the system. These strips vibrate and audibly alert driver that something in the system is different.

Rumble strips assume that the driver will make a mistake. Rumble strips are designed to work when the driver screws up. In fact, rumble strips will *not* work when the driver is actually behaving as road designers assume drivers should behave. Rumble strips only work when the driver becomes inattentive.

What rumble strips do is allow the driver to have a reasonable warning that something is wrong. The driver may then correct their behavior before an accident happens. The screw up is still there, the only difference is that the driver has time to recover. This example illustrates the importance of this idea.

Workers (and drivers) are at their best when they are allowed to detect and correct their performance in normal systems. Normal systems are most reliable

2
Why Think about Failure at All, Let Alone Think "Differently" about Failure?

QUESTIONS ABOUT YOUR ORGANIZATION:

- Is your organization as safe as it should be?
- How does your organization measure safety success?
- Is the production of product more valued than protection of people, really?
- Look around and see what is valued with money, time, and recognition—what is most relevant to your management team?
- Are your processes written to make work happen successfully, or to avoid compliance failure?
- Do you learn as an organization from your successes as much as, or more than, from your failures?

Here's what we have learned about safety...not everything, but some crucial things that seem to make a difference at work. If you will allow me the pleasure, I am going to make a list and, in turn throughout this book, discuss these ideas in more detail:

- The safest workers seem to be workers actually performing work.
- Workers are as safe as they need to be, without being overly safe, in order to get work done.
- Workers consistently create safety in practice while they do their work.
- Workers are constantly detecting and correcting variations in the work, the work environment, themselves, and others—to create safety.
- The work the workers do in an organization is remarkably different from the work that was planned for the workers to follow.
- Workers constantly have to make the work match the process, not the processes match the work. Workers are not empowered to "officially" change your process—only you have the power to change the process.
- Your worksites are undoubtedly a "hidden laboratory of alternative choices," in which we are always hoping that the workers make the right choice—and normally they do...until they don't.

- Your job sites, work areas, and facilities are normally exceedingly safe and dependable—worse things *can* happen than *will* happen—and because work is normally safe, workers are less careful.

In the opening case study, the Propane versus Freon tank failure, the failure is so damn stupid that it is hard to believe there would be any other organizational reaction to this failure than terminating the worker. The worker in this case study had it all—bad attitude, lack of ability to get along with his coworkers, "stupid" written on his forehead; he probably had bad habits and smelled funny. And in many ways this is the precise problem we are trying our hardest to combat. It is faster, cheaper, and a lot easier simply to claim the problem belongs to the worker and "fix" the worker. By simply choosing to purchase this book, you are telling me that you already have thought about why "fixing" people is so remarkably risky.

When you think of Human Performance, think of the three layers of your organization that are always present, during both failure and success. All three layers must be considered:

- The individual worker
- The organizational system
- The performance expectations set up by management

Too often organizations tend only to consider the first third of this list: the worker. The worker is where we put most of the effort; the worker is what we try to fix. The problem is if you only try to fix the worker, you are only getting ⅓ of the total organizational landscape. You therefore miss ⅔ of the potential problem: the system and the management expectations. You must teach yourselves to look at all three levels.

If the organization in the case study had fired the worker, it is a decent bet that in this case the problem would have gone away. The idea that there could ever be another perfect combination of a storage area, a tank smasher, that worker, that supervisor, that day, and that particular organization is incredibly low. But it is still there.

Is a "chance that event will happen is extremely low" an operational limit you can live with in your organization? Could your organization survive two failures involving Freon tanks being smashed dangerously?

By understanding the entire story, the whole system that "housed" the event, you can begin to understand not only the failure in the case study. You can also understand the other failures that are waiting to happen within that system, that organization, even that same work team. The ability to investigate an accident before it happens seems to be an extremely positive step forward in the understanding of not just failure, but the systems, people, technology, and components that must be present for failure to happen.

THERE ARE LOTS OF WAYS TO BE SAFE, BUT FAILURE SEEMS TO HAVE A SINGLE PATH

This part gets a little tricky. My advice is not to over-think these next ideas, but to understand how vital it is to your future success to recognize that these ideas have power over how you react and respond to failure. In short, the failures you look at in your organization could have only happened the way that they happened. If anything else had happened during the event, the failure would not have taken place. This idea is always true in retrospect, and seems to need the addition of the word "duh!" every time it is said in a meeting or event review.

But ponder this idea for a moment. Knowing the nature of all the things that had to have taken place in order for failure to happen is pretty valuable knowledge about not only the failure event, but also the organization that hosted the failure.

> How much of what happened always happens?
> *Answer*: Almost all of what happened always happens.
>
> How much of what happened has never happened before?
> *Answer*: Usually nothing different is recognized until after the accident happens. Remember this: it is exceedingly hard to notice something that doesn't happen.
>
> How relevant is it to know the difference between the two sides of the Freon/Propane story: the first story and the second story?
> *Answer*: Knowing both sides of the story is the most vital part of understanding the event. In many ways our entire job is to be able to look at both sides of normal and abnormal factors, and make sense of this information for the organization.

We are often lulled into believing that if we can just find the single, root place where the world went wrong we will have done our job and will then be able to make work a safer place for all. The problem with this idea is that we are learning that there is no one, single root cause that is the problem. The problem is always the relationships, causalities, or spaces around that single root cause, which seem to be in a constant state of motion and variability.

Knowing what was normal, what was not normal, and the entire gradient between normal and not normal begins to tell a different story of the accident, incident, failure, error, mistake, or upset.

WHAT IS FAILURE?

It is pretty essential to get our vocabulary together. So let's start with the understanding that when we use the word "failure" we are actually meaning the universe of words we can use to describe a deviation from expected outcome—DFEO. Failure is a good word to use because it is a large word that encompasses many other ideas. Accident, event, mistake, operational upset, flaw, deviation from the norm, screw up, foul up—these are all words that fit within the notion of failure. So when you see or say the word failure you are opening up all of the above mentioned situations in a quick and efficient manner.

Failure as a term is also significant because it crosses professional boundaries well. Quality professionals understand failure, production professionals understand failure, budget professionals understand failure, and, therefore, you don't have to spend much time translating these ideas internally.

Perhaps as significant, failure as a word translates well between professions and between cultures. There actually is not a direct translation for the word "safety" in Spanish, Japanese, or many Middle Eastern languages (all cultures I have spent much time working with Human Performance issues with in the last couple of years). But the word failure moves easily between concepts and languages in many cultures.

Why some cultures don't have a word for "safety" and do have a word for "failure" is, to a great extent, a real fundamental part of this book, and is pretty darn interesting. Safety, in many ways, and in many parts of the world, is a luxury that exists when the rest of societal needs are met, or close to being met. It hardly matters if the workplace is safe if you are more prone to being killed at home or on the streets. The luxury of safety for those of us in the developed world is no luxury at all—it has become a necessity and now a true business advantage.

Let's further define failure using Eric Hollnagel's definition. Hollnagel, an expert in this field, says, "A failure is the unexpected combination of normal performance variability." Hollnagel goes on to discuss that accidents don't happen because workers in the field have gambled their lives, and the organization's reputation; failure happens because of three key ideas.

First, failure happens because the worker believes that what is about to happen to them is simply not possible. This is so profoundly true that it seems simple and obvious. However, after a failure takes place we so easily move to a place where we see the entire picture. We know how the story started, what happened in the middle, and how the story ends, so we can see every possibility for the failure along the way. The worker has none of that information, and, therefore, does not have the same story or view of the failure within the event that we hold after the event.

Every time you drive home from work you fall into this potential failure mode. I will bet you that you don't perform a major safety systems check, take a pre-drive walk around your car, or even check your mirrors before you go home. I will bet you hop in the car, and drive home. That is what most of us do every single night of the work year. Why? Don't you know better?

Sure, you know better, but you also know that that same car got you to work in the morning without failure, so why would it not get you home that night? You have no indicator that anything is wrong, or even could be wrong. So you just hop in the car and go. Workers do the same thing at work. It worked last time. It worked the last 10,000 times. Normally, it always works OK. Why would it not work the next time?

All of us know that success one time does not lead to success the next time. Yet, success one time actually does usually lead to success the next time. In fact, failure is extremely rare. Your car does not break down very often. In fact, your car is much more dependable than it is not dependable. Think about this idea for a moment. Your car runs many, many more times than it doesn't run. Your car is more successful than the occasional mechanical failure. The same goes for workers.

Secondly, failure happens because failure often has nothing to do with the tasks and processes that the worker is currently doing in the field. We focus exceptionally well as humans on one thing at a time. Current neuroscience is telling us that humans are not particularly proficient at multitasking. In fact, neuroscience studies show that humans don't multitask—instead, humans share attention, some of us are fast attention sharers, and some of us are slow attention sharers. Depending on the context, we often find ourselves in situations where we cannot share attention at all. Think of the moment you pull a rental car out of some foreign airport and have no clue which way to go. At that moment, you are focused entirely on one thing—driving in a direction that will get you to your destination. You stop attention sharing, you turn down the radio, and you are not talking on the phone. You are driving for your life, or so it seems.

That same notion is true of things like texting and driving, or welding and catching your clothes on fire. These things are connected, but they are usually connected post-failure and in such a way that the connection becomes obvious and simple—after you see the connection.

Thirdly, failure happens when the worker feels the possibility of getting the intended outcome is well worth whatever risk is present in the work environment. The worker makes a trade-off between being safe and being productive. You want productive workers. The worker chooses to take a chance and guess that this iteration of this process will not fail this time and in this way. In other words, the worker looks around to see if any manager is watching and goes for it. The pop culture reference to this is "get 'er done!"

Many of you are fully prepared to argue that that is gambling. It may appear to have all the components of gambling. It may look like gambling. Hell, it probably smells like a gamble. Problem is, it's not gambling at all.

When you think of worker intent, or, as we have discussed it so far, gambling workers, think of how serious the outcome is to the notion of risk. In fact, think of this by using one simple word: surprise!

If a worker looks around and takes a chance and it fails, and the worker is *not* surprised, then the worker knew better than to take the chance, and probably was really gambling with your organization's reputation. But, if the worker looks around takes a chance and it fails, and the worker is genuinely surprised by the failure, then the worker was trying their best to get a difficult task done. It honestly is that simple, and is the greatest test of intent that I know personally. What do I normally discover when I am sent in to investigations? The surprises. I find many places in the middle of failure contexts that are just plainly and simply a surprise to the worker. I use the word surprise in my write-ups of events. I find the word and the idea of a surprise condition that the worker discovers while doing work to be a concept that is well understood by the people doing the work, and also well understood by the people who manage the people who do the work.

TYPES OF FAILURE

There are primarily two types of failure. These two types of failure are noteworthy because they help us understand how to move forward with the process of understanding not only how to investigate a failure, but also how to create corrective actions that will ensure that this failure, and others like it, is less prone to take place in our organizations. Knowing the type of failure allows us the power to know how to proceed in organizational understanding and learning. Soon we will discuss the power of learning in creating safe and efficient workplace systems.

- **Individual Failure**: A failure happens and the worker is not protected from the dangers present in the work environment. The consequence of the failure is almost always to the worker or workers in the event. This category of failure would include events similar to cuts, slips, trips, and falls; chemical or hazard exposure; strains. In general, failing to protect the worker from anything physical happening to that worker is an individual failure. Hurting people is an awful thing. Hurting a worker is seen as a failure of the worker not to operate safely.
- **Organizational Failure/System Failure**: A failure where the organizational systems allow some type of threat to the system to have a consequence where many people are adversely affected. In a system failure, someone or something has been able to break through the many layers of defenses that were thought to be in place protecting the facility, its people, and reputation. What you often find is the profound difference between thinking there are many layers of defenses, and actually having the many layers of defenses in place. Organizational events are things like oil spills, reactor coolant failure, single point recycling storage pens for collocating different small gas containers, or even air traffic controllers sleeping during the nighttime hours when no planes are scheduled to land at an airport.

Why do you care about these two failure types—and I might add you certainly do care about these two failure types? Because failure is almost always much more complex than it may at first appear.

Dekker says it something like this: it is much more seductive to see Human Performance as puzzling and perplexing, as opposed to complex. We become much more interested in the "why" question and somehow miss the "how" question about failure. Where we see puzzling "whys," the story of the failure often tells us something quite different. The story of the failure finds that success is complex: safety critical work depends on expert human performers performing work in an expert way. Our organizational systems tend to run degraded (from the launch of a system until it runs to failure), and plans, processes, rules, regulations, and procedures that we assume guide work towards safe operations are almost always incomplete.

You use the types of failure—something you have done automatically your whole career to understand where to start and stop the telling of the story of the failure. Most of the stories I tell in my work are stories about the way the organization failed the worker, and almost never about how the worker failed the organization.

The concept gets clearer when you break a failure down a bit.

> Caution: I am not asking you to deconstruct the failure down to its smallest parts. Near as I can tell, examining an event by taking the event down to its smallest parts has little value at all for you, for the workers involved, or for the organization. What we want to deconstruct is not the failure, but the way you and your organization understand the failure.

This next idea will help a lot.

FAILURE HAS THREE PARTS

Every failure that takes place in your workplace can be divided into three distinct parts. Let me rephrase that statement: failure must be divided into

three parts in order for you and your organization to begin to understand what is happening. You must be able to understand how the failure environment changed while the actual failure was happening.

The three parts of a failure are straightforward, and immensely important:

Part 1. The Context: Everything that led up to the actual failure event

We will call this "the context of the failure." Sidney Dekker, my friend and challenger on all things involving my organization, calls this part of a failure "getting into the tunnel with the worker." Everything—and I do mean everything—that is happening while the work is being done plays a role in the outcome of that specific work. We must try to capture and understand all the things that are happening to set up this failure to fail. The problem is that we must look both at the timeline of the event (a rather linear view of the world and not at all truly representative of reality), while at the same time look at the complex relationships that exist between all the many moving parts that had to align to cause this failure to actually be triggered by these workers.

Part 2. The Consequence: The failure itself

Someone or something triggers an unexpected outcome that for our discussion is seen by the organization as a bad outcome. Oddly enough, this is both the easiest part to identify and the most uninteresting part of the failure. This actual failure shows itself to us in many ways. All of these ways will be a clear deviation from an expected outcome. Often these ways are tragic and horrible. But the failure itself is only significant in the sense that it is the event. The failure will have an immediate consequence and will drive an often immediate organizational reaction.

Part 3. The Retrospective Understanding: Everything that happens after the failure happens

Call this the organizational reaction. Ah, the power of retrospective. This is where understanding the three distinct and different parts of a failure presents its best and most valuable pay-off. Everything that happens after the failure is colored with one giant rational advantage—that is, the fact that from this point on you and your organization know how this story ended. Remember, those workers back up in Part 1 of the failure had no inkling of this knowledge. Had they the rational advantage of knowing this failure would happen, it is a pretty solid bet that they would have done something—anything—differently in order to avoid the Part 2 of this event.

All three of these parts of a failure *always* exist in every failure, no matter the size. All three must be recognized and understood. Perhaps most importantly, you and your organization must be able to understand that all three of these failure parts are different, and must be defined and understood as different in order to understand fully the "story" of the failure itself.

WHEN UNDERSTANDING FAILURE, SOME THINGS ALWAYS MATTER

All divisions of failure aside, there are a set of core factors that seem to be present and relevant in all events. These core factors are necessary and collectively sufficient to help you tell the story of the failure. Without the core context information, your story is incomplete.

Core Context Information List:

I must have an explanation of the failure

I don't need the causal analysis document; I don't want the fault tree; I am not even sure I need a timeline; but I have to have an explanation of what happened. I am surprised by how often I go to the field and look at amazing report documents that don't thoroughly explain how the event happened. I actually look for a story that can be told of how the failure happened. Remember your junior high school English class? A story has a beginning, middle, and an ending; the story must make sense (have fidelity), and move through time logically and sequentially, but not necessarily linearly.

I need to understand how the energy of the failure was or was not isolated from the worker or workers involved in the failure

That sentence is written in High Reliability-ese; let's see if we can reframe this idea in a way that makes sense to people who are in the field. I need to know what went wrong and what went right. Where did the failure start, and when did our systems kick in to mitigate further failure (if at all)? Which all leads to the third requirement...

I need an understanding of how some type of defense or defenses built into the organizational system were either not present, or if present for some reason failed

This chapter of this book seems to contain many absolutes. Here is one more: every failure that takes place happens because some defense either did not work or was not present in the work environment. This defense comment is always true; albeit incredibly retrospective, it is always true. Think of it like this: I am not surprised that air traffic controllers, working alone in the middle of the night at an airport with no scheduled landings or takeoffs, take naps. I

would be surprised if there were no way to wake the controllers up when they were needed to perform their mission. A horn or a buzzer connected to the radio would be an excellent defense against an exceptionally severe lack of stimulation and an incredibly high amount of boredom.

I must be ever vigilant in remembering, "size does not matter"
The consequence of a failure does not determine the importance of the failure. We are humans, our organizations are made up of humans, and we all believe that consequence drives outcomes. I have spent the last 15 years of my career asking this question of classrooms full of people just like us: "What is the difference between a big event and a small event?" The answer, and I think I have heard them all, really should always be, "nothing." All events allow us the opportunity to learn. Small events, those events with small consequences, are often much richer in context and story than what our organizations define as large events.

I must always remember that bad things don't always happen to bad people
Or at least stop thinking the person is bad (or allowing the person to tell me that they were for some reason bad). The truth of the matter is that bad things happen all the time to very good people, very experienced people who had no intention other than doing the very best job they could do for their organization. As silly as this seems, not thinking people involved in an event are bad is difficult and will take deliberate attempts on your part to not label the workers involved as something less than perfect.

Here is what you know from years of doing safety work. Sometimes big, horrific events are actually quite simple, straightforward failures. A bolt snaps, scaffoldings fall, and the consequences are terrible. The other side of that coin is that sometimes tiny events have profound learning implications for the organization. A small event happens, and you suddenly realize your entire system is flawed and must be immediately strengthened. Event size is not nearly as powerful as event learning. Don't let the drama of an event (or lack of drama of an event) dictate how much or how deeply you respond.

The problem is that the only way you know which events are information rich is to look at them. Therefore, every event demands event learning. Every event must demand event learning.

Case Study
The *Titanic*: A Story to Help You Rethink How You Think About Failure

Before you read this case study, a study you will undoubtedly know something about, ask these three questions as a pretest of your failure understanding ability:

- Whose fault was the sinking of the *Titanic*?
- What caused the *Titanic* to sink?
- What corrective actions would you have recommended for the next ocean liner to be built?

The *Titanic* was built to be the flagship for the White Star Line. It was the largest and most opulent passenger ship in the world. This stately ship was "state of the art" for its time. White Star designers touted a ship built with all of the luxuries of the "gilded age," with countless fineries and undeniable craftsmanship: a Turkish bath, teak wood furniture (even in third class), a heated pool and lending library. There were crystal chandeliers, hand carved banisters, and gold leaf detailed throughout the ship. No expense was spared in its luxury and beauty.

The same eye for detail was brought to the technological design of the *Titanic*. New technologies were integrated. It boasted a 5 kilowatt Marconi wireless radio (invented just 11 years before), a telephone system, two bronze 10′ propellers, and a double plated 2′ steel hull. The builders and designers had believed they had built the safest ship in the world. In a word, it was "unsinkable."

THE SAFEST VESSEL OF ITS TIME

The *Titanic* was to be the safest vessel of its time. A major improvement to its design was its hull. The hull was subdivided into watertight compartments. This was a new safety design and had never been done on any other passenger vessel.

The ship was monumental, measuring 175' high, with 9 decks. At its top speed it could travel 23 knots (26 mph). Fully loaded it had a passenger capacity of 3247 in all three classes, and a crew capacity of 885.

The *Titanic* began its maiden trip from Southampton, England to New York City on April 10, 1912 with 2,223 crew and passengers. The first class passenger list boasted the "jet set" of its time. The list included Benjamin Guggenheim, John Jacob Astor IV and his wife Madeleine, and Isidor and Ida Strauss, the owners of Macy's department store. In the second and third classes were many families moving to start a new life in America.

SYSTEMS ARE NORMALLY DEPENDABLE...

The first few days of the journey went incredibly smoothly, leading the ship's owner, Bruce Ismay, and its experienced Captain, Edward John Smith, to begin to feel a bit more at ease. They considered pushing the ship's speed a bit. The normally choppy North Atlantic seas were calm, and the weather was cooperative. Every hour the crew would drop a bucket over the side of the ship and bring it up, drop in a thermometer, and measure the temperature of the water. The temperature would indicate the presence of ice. The crew were rotating their shifts in the eagle's nest at the top of the ship, spying for ice. The ship was also receiving Marconi messages from other ships in the area providing the location of ice fields moving into the shipping lanes. These were the standards of practice at the time. The crew was making sure that they, and more importantly, the passengers, would have a smooth trip to New York.

UNTIL THEY AREN'T

In the late evening of April 14, 1912, right after the first class passengers enjoyed a magnificent 10 course meal, at 11:40 pm, the *Titanic* hit an iceberg. Only seconds after the crew saw the ice, the ship was hit on its starboard (right) side, with the impact lasting a few seconds, leaving a giant gash in the steel hull.

In just under three hours, this dazzling, beautiful liner sank to the bottom of the North Atlantic, resulting in the tragic demises of 1,500 people. This is the deadliest maritime disaster in our peacetime history.

HOW?

How did this happen? This ship was "unsinkable." There were systems in place to ensure everyone's safety. The *Titanic* was a technological marvel. It had

everything to conquer the predictable hazards of transatlantic passage at the time. Multiple safety systems existed around every analyzed hazard that was known. Safety was not just a significant value; safety was also a marketing strategy. It would be hard to find evidence that reliable and safe operations were not one of the most important operational values. Nothing could have or should have ever happened. The entire mission of the *Titanic* was not to sink, and yet it did.

Let's continue the story of this complex, multilayered failure. All of the designed safety systems, and trained and safe workers were in place, and functioning exactly the way they should have been. Nothing unusual happened; in fact, what happened was an amazing and unexpected combination of many normal things.

We will never know the exact size of the iceberg; however, early newspaper reports of the time estimate that the iceberg was around 50 to 100' high and 200 to 400' long. However, an interesting point about the size of icebergs is that only a teeny, tiny part of the berg is visible above water; the greater part of the iceberg is underwater, and is 10 times larger than the visible portion. Most icebergs expose about 1/10 of their mass above water, leaving the other 9/10 of their mass below the water line. In this part of the Atlantic, icebergs were an expected, accepted, and normal part of sea travel.

Sixteen separate cells of *Titanic*'s hull could each be isolated and closed with special doors in the event of an emergency. This was an amazing new design for any passenger ship. However, to be watertight they would have to be able to be closed like a box on six sides, making parts of the ship impassable to crew and passengers. This unintended impediment was operationally impractical, so doors were placed in the bulkheads to allow access. Because of these doors, an adaption in the ship's design so that the ship could be operated, the watertight bulkheads were not really watertight. Because of the design confidence in these separate cells, no consideration was given to the idea of what would happen if water did penetrate them. The cell walls were tall but not watertight all the way up to the ceilings of the ship. Water could flow over into each hull compartment. All of these new and highly technical defenses were designed with the best intention of safety. This system was the most modern and the safest ever developed.

In the crow's nest of the ship, the place high above the ship where crew kept a lookout for ice, there were supposed to be a pair of binoculars. White Star leadership made a relatively normal last minute change in the equipment officer. The old equipment officer forgetfully took the key to the equipment storage locker with him when he was transferred from the ship. *Titanic*'s new equipment officer couldn't find a replacement key, and because there was no access to the equipment storage area, the lookouts were not issued with any binoculars.

The White Star Line's official policy was that binoculars were supposed to be given to all lookouts. Since they only had a few binoculars on the ship due to the error, these binoculars were given to the highest ranking officers, not the lookouts. On the night of the collision, the weather was cooperating. It is very difficult to notice something that doesn't happen, and in this case in this sea condition, nothing was happening. In these conditions it was very difficult to see the iceberg. Churning or violent water would have caused breakers, with some flotsam and jetsam, around the base of an iceberg in the water, making it easier to see it from a greater distance.

Titanic did not strike the "typical" iceberg. It wasn't white, gleaming out of the ocean like most icebergs. Witnesses say that it was almost invisible from continuous melting and refreezing. This clear surface of the berg reflected the water and dark, clear night sky like a mirror, thereby making it almost black, or the same color as the evening ocean. It would have been the ocean equivalent to the black ice seen on roads. This type of iceberg is nearly impossible for a qualified and competent lookout to see in these conditions, even with the keenest eye and the highest professional motivations. The very best lookout could and did easily miss this hazard. Not because the lookout could have done it better, but because this hazard was so difficult to identify in that condition.

The *Titanic* crew had a total of 6 hours for sea trials before they sailed from Southampton to New York. The sea trials were pretty much a formality. The members of the crew who stoked the fires were on board, and representatives of the White Star Line, but no domestic staff were aboard. Once the surveyor from the British Board of Trade, who was also present, saw that everything worked to his satisfaction, and that the ship was seaworthy and crew were fit to service passengers, he signed an "Agreement and Account of Voyages and Crew." This document was valid for 12 months, and that deemed the ship seaworthy. They did have additional time to prepare when the ship traveled from the shipyard in Belfast, Ireland to Southampton. The crew had 6 hours of practice to prepare for its first intercontinental cruise. Many of the crew did not know where stairwells and exits were located, especially in the second and third class areas. Some of the exits were locked, keeping the passengers below deck. Familiarity with the ship was clearly as complete as it needed to be without being overly complete. This short sea trial was perfectly acceptable for the most technically advanced and safest ship ever built. This sea trial only becomes incomplete, rushed, and wrong because the *Titanic* was lost at sea.

SAFETY SYSTEMS COMPLIANT

The *Titanic* was fitted with a total of 20 lifeboats, comprising three different types. Initial designs for the ship had included 48 lifeboats. However, in an uncharacteristic but financially reasonable cost-cutting measure they cut 28,

leaving only enough lifeboats to hold 58 percent of the passengers and crew. Because of the added safety design features built into the ship itself, it was acceptable to determine that the additional lifeboats were not necessary. The White Star line was fully compliant, well within the law. The British Board of Trade regulations stated that all vessels over 10,000 tons must carry 16 lifeboats with a capacity of 5,500. These regulations were becoming antiquated since the sizes of passenger vessels were growing enormously. The *Titanic* weighed in at 46,000 tons. Not only had the White Star Line fulfilled their legal and regulatory commitments, they were providing more lifeboat accommodations than were required. The high casualty rate in the sinking was due to the fact that although White Star were meeting the current regulatory requirements, the Titanic only carried enough lifeboats for 1,178 out of a passenger and crew list of 2,223. An outcome from the investigations into the disaster in both England and the United States was a drastic change in the antiquated regulations.

Recent tests on the steel from the *Titanic* reveal that it was far more brittle than the steel used on modern ships, but it was the best available at the time. It is easy to look back and say that they used inferior steel. However, that is the advantage of time, retrospect, and knowing the ending of this tragic story. The White Star Line spared no expense in this part of the construction, and used what was the best material of the time.

PROFIT OVER PROTECTION

Titanic housed an office for the wireless company The Marconi Company. The wireless allowed the passengers to send and receive cables, a luxury in communications at that time. It also allowed other ships in the area to transmit the coordinates of ice flows. Policies on these types of cables were simple: the cables with the locations of ice fields or icebergs were given to the bridge officer, who would then tack these paper cable transcriptions on to a board in the bridge. This provided current information to the officers or those coming on the next shift.

The Marconi officer working the night of the sinking logged that he gave a cable to Capt. Smith in the early evening. He remembered that the Captain was in a rush and that he placed it into the right-hand pocket of his jacket with just a precursory glance. There were six iceberg warnings received by the *Titanic* on the day of the disaster. One message got to the Captain, but the others were ignored by the wireless operator who was busy transmitting passenger messages. Passenger messages were the Marconi's bread and butter. The wireless was clearly less available for operational messages (which generated no money for the company) because of the substantial profitability of the passenger use.

NEW RULES OR OLD RULES?

Company policy only required the wireless offices to be operating until midnight. *Titanic* hit the iceberg at 11:40. Most of the ships in the area had turned in for the night by the time the ship was transmitting calls for help. For several years the distress call for ships was CQD, which translates as: CQ = attention all stations, D = distress or danger. A short time before the *Titanic*'s launch, an international convention introduced a new distress call to supersede the traditional CQD. The letters chosen were SOS, not because they stood for anything in particular (although rumored to be Save our Ship), but because they were simple enough for the new amateurs flocking to the wireless craze to send and receive. *Titanic* was tapping out at 12:15 "CQD MYG" (the ship's call letters). *Titanic* was the first ship to have ever used the newly negotiated distress call SOS. The ship's wireless crew tapped both CQD and SOS in order to cover all bases.

At 12:19, on the Cunard Line's *Carpathia*, just as the wireless operator prepared to turn in that night, he heard the call. The ship was less than 60 miles away. The *Carpathia* picked up and saved the *Titanic* survivors. Had the operator turned off his wireless, the distress call would not have been heard and no rescue craft would have responded.

At around 10:30 pm, the liner the *Californian*, a cargo ship, had stopped for the night on the edge of an ice field. As was the custom of the time, the wireless operator had turned the radio off and had gone to bed. The ship's night crew noticed a behemoth liner stop an estimated 6 miles to the south at 11:40 pm.

Shortly after midnight, the Captain of the *Californian* was informed by the crew that an enormous passenger liner was shooting rockets into the sky. The crew concluded that this ship had anchored for the night and that there was some type of party. These were not fireworks but *emergency* flares. At 2:20 am, it was noticed on the *Californian* that the *Titanic* had disappeared. Those on the Californian believed the other ship had steamed away. The *Californian*'s wireless operator woke up in the early morning and turned on his Marconi radio. It was at that moment the crew learned of the tragic fate of the *Titanic*. In both the British and U.S. Senate inquiries into the disaster, the Captain of the *Californian* insisted his ship had been many miles north of the *Titanic*, and that he could not have reached it in time to rescue passengers. However, many survivors testified in the investigations that they had seen the lights of another ship roughly six miles north of the *Titanic*. Both international inquiries concluded that the *Californian* may have been only six miles or so to the north of the *Titanic*. The *Californian* could have reached the *Titanic* before it sank.

DISCUSSION

Since you have now read this case study, again ask these three questions:

- Whose fault was the sinking of the *Titanic*?
- What caused the *Titanic* to sink?
- What corrective actions would you have recommended for the next ocean liner to be built?

Remember these questions and your initial answers. Have you reinforced your initial answers, or are you now reconsidering your initial answers?

The tragic story of the *Titanic* is a story about so much more than an iceberg. It is easy and profoundly wrong to assume that an iceberg sank this ship. It is also profoundly dangerous to think that an iceberg caused this event, because focusing only on the iceberg takes the safety learning attention and the allotted corrective action resources away from the real weaknesses and potential learning opportunities of this event. Worse yet, blaming the lookout or the captain, although a tight story and emotionally satisfying, is not terribly accurate. No one person alone was able to have this failure: this failure was a product not of the crew, but of the relationships between the crew and leadership, the crew and passengers, and the designers of an unsinkable vessel.

No one person or thing caused this failure. Instead, a whole series of small, seemingly innocent, normal activities combined in such a way as to cause the safest vessel ever built to cruise the seas to sink. No mechanical reason alone caused the ship to sink. However, the mechanical watertight cell walls not going all the way up to the above boat deck is, in retrospect, a failure. However, had there not been doors in the cells, water would not have been able to move from cell to cell. If there had been no doors the crew and passengers could not have moved throughout the ship during normal cruising.

Interestingly enough, no one error caused this failure either. In fact, there is an argument to be made that in almost all cases, no errors were made in context. The errors only became apparent after the *Titanic* sank. Confusing emergency flares for party fireworks was not an error until the realization was made that the ship that was six miles away having a party was actually sinking. Even the captain placing the ice flow Marconi transcription in his pocket would have had little effect if the other five transcriptions had been completed.

All of these small, unpredictable components of this failure seem very important and very dependent upon one another when you think of how the *Titanic* failed; however, in the context of operating an unsinkable, technologically advanced ocean liner on her maiden cruise, no one event would have come close to triggering even a low level alert. It is not until after the failure happens that these seemingly normal deviations from expected outcomes become very important warning signs.

This is a complex (many parts, many people, many systems, many stakeholders, many goals, many pressures—all closely related to one another), adaptive (decisions being made in real time based upon information that is readily at hand in order to create the best outcome at the time) failure. No one cause. No one person to blame. No errors to discipline or train away.

3
Change the Way Your Organization Reacts to Failure

QUESTION ABOUT YOUR ORGANIZATION:

How does your organization act after something unexpected happens?

1. We have an immediate meeting and decide what happened, and how to fix it.
2. We pull the worker(s) involved into a meeting about the event.
3. We name, blame, shame, and retrain.
4. Somebody's going to get a phone call—probably me.
5. We see it as a chance to learn and probably will compile a team to learn as much as we can.

A STORY

This is a story about a maintenance man, and a small carnival operation headquartered in the Midwestern part of the United States. The carnival traveled from town to town, county fair to county fair, with 20 amusement rides, 6 dark rides, and several games of both skill and chance on the Midway.

The equipment in this carnival was aging, and demanded almost constant attention from management, ride operators, and the maintenance man. There was enormous pressure, because of the small, mom-and-pop size of the carnival, to keep every amusement ride in running order. After all, a broken ride cost the carnival company money, and did not make the carnival company money. In order for the maintenance man and the rest of the carnival workers to get paid on Fridays, all the rides needed to run as full as the can all week.

Times were hard. It seems that not as many small towns had reasons to have carnivals, and bigger carnival companies with better, more exciting rides were picking up the state and county fair jobs. The pressure to find work was high. The pressure to delight the customers became more intense. The problem was that while performance pressure was increasing, resources to ensure the expected level of performance were declining. Times were hard and seemed to be getting harder.

Several weeks previously, "The Whip" had started making a strange grinding noise. "The Whip" was a high speed swing ride, where the riders sat in small cars suspended by cables from a central spinning motor axle. The ride was designed to start slowly, then turn to the right in a circular motion, eventually turning faster and faster. When the speed neared its maximum rate, and due to the nature of centrifugal force, the little cars filled with riders would be traveling at about 20', parallel with the ground, for about 4 minutes.

To say that "The Whip" was exciting is probably an understatement. "The Whip" was fast, fun, and dangerous. Everything you want from an amusement ride. But the ride started to make a grinding sound; the sound was an indicator that something was not right with this ride. Shutting it down was not the best choice. The ride was one of the carnival's more visually appealing rides. The carnival company had evidence that something was wrong with the ride, and had to make a safety decision with operational consequences.

Anyone could stop the ride if they felt it was unsafe, but a little grinding was not unusual, and the line of riders had been long and steady all weekend. "The Whip" could be stopped all right if the grinding got much worse—the problem was there was no standard for what worse grinding sounded like to each operator.

So the decision was made to perform maintenance on this ride on Saturday night. As soon as the fair closed, the maintenance man would troubleshoot the ride, find the problem, and fix it. With any luck, the entire ride would be up and running and full of riders by 10:00 am Sunday morning, when the fair opened for its final and most profitable day.

The maintenance man tore the ride down, and checked everything he could think of checking in order to fix the ride. The more he worked on the ride the better the ride sounded. It was almost back to working order, and seemed to be operating at about 90 percent. The maintenance man did not have time in one night to take apart the transmission housing, and check the main linkage between the transmission and the ride arm axle. That process was easily a two-day job. He would hold that task until Monday and Tuesday. The beginning of the week was the time when the carnival either moved, or in rare cases gave their workers a couple of days off. This partial repair would have to do—it wasn't perfect, but it was good enough to make it till Monday.

At about 6:30 am Sunday morning, the maintenance man tightened his final bolt, and turned the ride to the operational setting. He checked to make sure the operator's logbook noted the repair, and did the paperwork to make the ride operational for that day's riders. The maintenance man then went back to his travel trailer and went straight to bed. He was getting too old for this crap, and he was certainly feeling it when he had to do an all-nighter. Sleep would be welcome, and come easy.

But it didn't, not at all. In fact, the maintenance man could not shut off his brain. He just kept thinking was "good enough" really good enough? Many

lives would depend on this ride not failing, and the maintenance man could not say for certain that the grinding was not a sign that the main coupling was getting close to failure. Had he greased it enough over the past couple of years? Was the grinding a sign that it was wearing badly and may be close to breaking?

On Sunday morning, at 9:15 am, the maintenance man got back out of bed, went to the management, and told them what he was thinking. He talked to his bosses about the potential of the coupling giving way and literally throwing riders across the fairground at 40 mph, at an altitude of 20' above grade. It was too risky, and he just did not feel right about his work. The carnival must cancel operations of the most exciting, highest grossing ride on the biggest day of the fair.

And so the management canceled the ride. The ride manufacturer was called, and sent a field inspection team to do an emergency inspection of the ride. The manufacturer's field inspection team went ahead and pulled the transmission housing off the ride, and discovered, to everyone's worst fears, that what was once a 2½ inch-thick steel axle had worn to less than a quarter inch. Had the ride operated that day, *that day*, 50 people would have been seriously hurt or killed, and the carnival company would have gone out of business.

In essence, the maintenance man, and his ability to gather the nerve and fortitude to call stop on "The Whip," saved the lives of the riders, his carnival company, and perhaps the carnival industry. The maintenance worker was given a $1000 bonus, recognized by both the county and the state for his heroic efforts, and featured in the local newspaper and even on the cover of the professional magazine of carnival operators.

All in all, the worker and the company together did exactly the right thing, and safe operations saved the day. The world was safe, reliable, and productive once again.

But, just for a moment, let's play out another ending…

What would have happened if at 9:15 am on that same Sunday morning, that same maintenance man had told that same story? The ride would have been shut down in the same way. The manufacturer would have been called in to the fairground. The manufacturer's field inspection team would have flown out at the carnival's expense to the fairground in that small Midwestern city. However, imagine that this time when the field inspection team took the housing off the transmission in order to inspect the coupling, all they found was nothing. The axle was perfectly greased, and the noise was simply normal. The grinding sound was well within the operational parameters of the ride's specification.

Would the maintenance man have been rewarded? Would he still have received the $1000 bonus? Would the county and state have recognized him? Would he have received any praise at all?

...Or would he have been seen as overly careful, foolish, and not a "smart" decision maker?

DISCUSSION

This is an interesting story. How you react to failure, or react even to the *potential* for failure, matters to how safe your operations are operating. It is easy to celebrate a worker who exercises good judgment with a questioning attitude when the worker's questions lead to a discovery of a potential failure before it becomes a failure with serious consequences.

It is not so easy to celebrate exactly the same behavior and decisions when the outcome is not a potential high stakes failure, and, in fact, costs the company a ton in lost revenues. The natural inclination is to say that the worker was too cautious, too careful, and overly vigilant— and all of those things are bad if there is, or was, no threat of failure.

And yet we are talking about exactly the same behavior, the same attitude, and the same worker who worries about a problem and reports the problem to their management. The decision to stop operations was a costly one, but when it avoided a real problem, a potentially horrible accident, the cost went from being a burden to a blessing.

In many ways, this story illustrates much about many organizations. You must first understand the many inputs that go into your reactions to events. It is not an easy task. It is sometimes costly, but ask yourself this: if you were the maintenance worker, and the story took the second ending—"nothing wrong here"—and you were "punished" for being overly cautious, would you ever report anything to your manager again?

Your biggest challenge is to create an organization that is not afraid to talk to each other about safety. This is no small challenge. To do this you must first understand that to expect trust, you must give trust. To expect communication and a questioning attitude, you must build an environment that fosters those behaviors. To do that, you must always think of how your organization reacts and responds (two remarkably different words and ideas) to failure.

WHY YOUR REACTION MATTERS

Something unexpected and bad happens...

...And everything that you and your organization does from this moment on will become necessary to creating everything else that happens—the critique or after action meeting, the investigation, the corrective actions, the whole

response picture. Time matters, because it is impossible to have an organization that is operating out of control.

This situation is tough, because your first reaction after an event has been to identify who made the mistake and think immediately about ways to fix that person. The old view of accidents was a pretty narrow view—simple, sweet, and cheap to fix. If you find the guilty person, and remove or retrain that person, you have clearly fixed the problem. Here is a problem with that old model: the "bad person" model tends to lead any organization down the linear, mechanistic path. The problem with this path is that it feels as if you have done something, have fixed the problem, when in reality that could not be farther from the truth.

It seems simple, almost too simple, but your organization's reaction to a failure tells your workers everything they need to know about your company, its managers, and what is going to happen. How you respond sets the stage for everything that will happen, and that has happened, and for how the event happened in the first place.

I am going to say something that you should question immediately. In fact, in normal circumstances I would never trust anyone that even came close to saying what I am about to say. If there ever was a time to call "B.S." on anyone, this would be that moment. Here it goes:

> I know how you can change your organization's entire safety, quality, security, or production performance for the better, immediately... Change the way your organization reacts to its failures.

It is a bet: if your organization is anything like my organization, when something bad happens your organization jumps almost immediately to some type of faultfinding posture. Who screwed up this time? The new philosophy I am proposing helps your organization move from the classic "crime and punishment" reaction to a "diagnose and treat" response to failure. That shift is exactly what we are hoping to create in the new organizational view of safety.

TOOL: BETTER QUESTIONS, BETTER ANSWERS

You must help your managers know this little but important fact. You must tell your managers that the way they react or respond to an operational failure matters—and it matters a lot. Coach, counsel, and press your leaders into knowing that they play an enormous role in event learning, and organizational culture.

All of this boils down to the premise that your mangers will need to know what to do. If you want them to do something differently, you must tell them

what that different thing is that they should be doing. Hence the following job aid. This is a card that we produced which asks the managers to be aware of both their reactions and responses, and suggests some better alternatives to ask when confronted with an event.

Nine Things Managers Should Ask First and in Order When They are Notified that Something Unexpected Happened:

1. Your response to an event matters!
2. Are the people OK?
3. Is the facility safe, secure, and stable?
4. Tell me the *story* of what happened?
5. What could have happened?
6. What factors led up to this event?
7. What worked well? What did not work?
8. Where else could this happen?
9. What else do I need to know about this event?

At my organization, we put these nine questions on a card and handed that card to our managers. I discussed this concept earlier in the book, but here is what my organization found. Our managers want to ask the right questions; we had not done a decent job of telling them better questions to ask. Interestingly, when our managers actively tried to gather explanations and stories, and not reasons and causes, we started to get decidedly different, much higher quality reporting of events.

With a better understanding came better responses and better solutions. With better solutions, we saw organizations drop not only the number of their events, but also the severity of events.

CRITIQUES, AFTER ACTION REVIEWS, FACT FINDING MEETINGS

> "This meeting is being called to understand how we had an accident in area 7 this afternoon. This meeting is not about blaming workers. We want to have an open conversation about this situation. Feel free to communicate openly and honestly to management...Now, tell me how you guys screwed this up this afternoon."

Actual Critique Opening

Your facility most likely holds some type of event review meeting after a failure. In larger facilities, this meeting can be quite serious. Smaller facilities may not

have the same formality and size of this type of meeting, but it is almost certain that this type of meeting happens in these facilities too.

These meetings are usually used to perform several critical tasks. The primary reason for these meetings is to understand what went wrong and decide what needs to be done, usually at that very moment—or moments after the meeting. The second reason to have these meetings is to categorize the event by some standard. Is the event in question a really an event at all? Is the event a product of bad worker choices? Is this event a result of management actions? Does this event fall someplace in between these two extremes?

But here is the problem with these after event meetings held after something terrible happens. These meetings become a little bit like a court hearing, and often not much of an opportunity to learn "how" the event or events transpired.

Here is a question you should ask about your organization:

> "When was the last time you had a full blown critique after a job was completed successfully?"

It is easy for your organization to find resources to understand and investigate a failure. We always have the time and money to do an investigation. In reality, because it is an event, we know that we will have to pay for some type of learning activity. The problem is that it is not easy or even practical to understand and investigate success. Success is normal. Success is common, and, therefore, we don't actually see a need to understand why a job goes well. Even if we did track successes with disciplined regularity, how would we find a data management system to keep all the success data that we would find?

Post-job reviews after operational success are an excellent opportunity that is often missed or wasted by organizations. These meetings ought to be an opportunity, open and honest, to tell the whole story of the success that just happened. Workers should feel that these meetings are fair and open, and that the management is genuinely interested in learning about all the factors that made up the success. Looking at failure with the same type of meeting you use to view success sends a strong message to workers and managers. The message is that your organization is interested in learning no matter the outcome.

Often the room itself is an indicator of what the meeting is like for the workers involved. You might be surprised to know that the management conference room is not the psychologically safest and most open place for the workers. Many times I have moved chairs and tables, brought in coffee and doughnuts, and asked senior managers to "non-attend" these meetings in order to create a space where the workers may speak openly and honestly.

If you are interested in building a culture in your organization that is able to report events in ways that best tell the whole story about a failure, you must be interested in changing this very formal after failure activity. It is helpful to be open to the reality that these meetings, which often feel truly open and safe to management, may give a completely different impression to the workers who are required to attend.

TOOL: MAKE THE FORMAL LESS FORMAL

Try these three things the next time you have an opportunity to have this type of meeting:

Don't just start the meeting telling the members of the meeting that this will be an open meeting, nothing will be used against the worker, and our goal is to learn—do something to demonstrate that this meeting is different. Work at doing whatever you need to do to make this learning opportunity feel and look more like a chance to gather information and less like a courtroom proceeding.

Make this meeting less formal, less "courtroom-like," and more about a chance to go about honestly learning about the story of the failure. Use the "9 questions managers should ask" when notified of an event in your organization to guide this post-event meeting.

Have a post-event meeting for work that was done successfully, and invite lead workers and opinion leaders, and ask them to tell you the story of what went right, and why we as an organization were successful in completing this task. Take this meeting seriously, and go into this activity with the idea that you want to learn about what works as much as you want to learn what went wrong.

Be aware that reactions to events tell workers everything they need to know about your company's safety program. If you do not attempt to "fix" this problem, everything, seriously, is communicated about how your organization will move forward. It is more important to focus on how your organization will learn from this event and avoid future occurrences.

Learning from success will surprise you. Successful completion of work in your organization includes just as much worker judgment and work improvisation, and just as many adaptive behaviors and surprises as work that fails. You can learn all the same stuff you learn after a failure, without the costs of the actual failure.

This whole process is exactly how we make safety better in our organizations.

TOOL: LEARNING TEAMS

One of the better tools my organization developed to increase and improve operational learning is a remarkably straightforward tool with a remarkably high level of effectiveness. This tool was developed to address the need we had to make event learning more effective for our facility. We knew there was lots of information on our job sites, but we were not getting much operational information. We had an immense challenge, and the challenge was this: we needed a way to engage our workers in event learning and a better way to get important information about our systems and processes.

We knew that we could achieve this if we could just tap into the informal learning that was taking place at the task level. We knew the workers always seemed to have a much more complete contextual understanding of how failures and near misses were happening in the field. We needed to capture this information in a way that reinforced and respected the workers' story, while at the same time was an effective communication tool with, and for, our management.

What we came up with was the learning team. A learning team is an ad hoc group of workers from any level within the organization that have been brought together to answer one question: Something has just happened—what should our organization learn from this event (or potential event)?

These teams, or, better said, this idea of being able to create a team at a moment's notice, could be called upon by a manager at any level of our organization for problems (or potential problems) that warranted a little bit more attention. However, this type of operational learning also needed to be legitimate enough to use as a part of event response for the regulator.

In your guidance to learning teams, you should tell the team members that you are not interested in "who" did what (we almost always know that information) or in "why" the event happened ("why" tends to lead to a more mechanistic failure picture); you are most interested in "how" the event transpired. You desire the story of the failure—with a clear beginning, a clear middle, and the ending. This can't be stressed enough: *how* is more important than *why*, and *why* is more important than *who*.

Once the team starts building the story, you ask them to produce some type of report or product. This could be a set of slides, a written report—whatever the team thinks is the best way to communicate what they have learned. You then have the team report this information out to some appropriate level of management.

If this all seems rather unstructured and loose, then that is because it is unstructured and loose by design. In your organization, you want a learning team to be a management tool that exists somewhere before a critique and well before an investigation. This tool exists not to create a formal record or some type of auditable document, but instead is a tool for the organization

to learn and correct from. You also want your learning teams to be a worker-owned process that provides a way to encourage worker ownership and to gather worker-level information about operations.

The challenge has been in trying to capture these learning team reports at the corporate level. I have found these teams incredibly valuable at the local level. It is my guess that this information will have potential value for the entire organization. At my organization, we are trying to crack that code even as I write this information for you.

The challenge persists that as this process is formalized, you will lose some of the honesty that you are getting by using this process. If this process becomes formalized and structured, the process itself will move from a way to gather real, leading safety information to a lagging deliverable to either the organization or to the regulators.

Try this method and see what you find. Customize this tool (and others) to fit your needs. If you communicate and manage these learning teams effectively, which is pretty easy—leave them alone and give them broad expectations ("tell us what we should learn from this event")—you will soon discover that learning teams become your corrective action planning program. Workers know what keeps them safe. Workers keep your organization safe.

Pre-Accident Investigation Tool

EXPANDED TOOL DISCUSSION

Have you ever watched one of the many CSI (Crime Scene Investigator) shows that have been on television recently? These shows tend to feature an almost "magically logical" deconstruction of a crime in order to track and capture the bad guy. It is done neatly in 60 minutes, and often involves what we have called in this book a more "mechanistic" investigation of the crime for most of the show, until sometime near the end of the hour the CSI team puts all the information together in a more holistic way in order to tell the story of the crime that happened. The story is usually a complete description of how the crime happened, what the motive was for the crime, and, most importantly for the CSI show, who is the culprit. The concentration is not only on who, but the why, what, when, where, and how of the crime.

In essence, the CSI shows are exceedingly powerful examples of the Human Performance view of understanding failure. The investigators strive to understand and explain what happened without judgment, in order to understand the story and to provide a just and honest conclusion to each case.

A just, honest, and concise explanation is precisely what every Human Performance investigation strives to provide. We look for the story of the story: a complete, context filled explanation of how the event happened is the best possible outcome of any formal method of understanding failure. Often the Human Performance investigation gives the organization information that is incredibly comprehensive, makes it easier to identify what to correct than with old school methods, and actually helps build a culture of trust and communication.

The only problem is that in order to get this CSI-type information you have to have a failure. There is nothing good about having an event in order to learn how not to have the very same event. Or is there? We are pretty sure there is.

We can often get the story of the failure before it becomes a failure. It is simply a matter of time, access, resources, and luck. Sometimes we discover problems in waiting. Other times we are surprised by events that we wished we had discovered beforehand.

Here is one way to help formalize and structure your prevention strategies by leveraging your organization's ability to learn. Think of your pre-accident investigation as a 6 part process.

1. Look for high consequence activities
2. Look for small signals that can indicate system weaknesses or problems within the normal work process
3. Look for error provoking systems steps and processes
4. Look for error likely conditions
5. Listen to your workers
6. Ask yourself what keeps you up at night?

Identifying Pre-Accidents in Your Organization

You can't fix everything. You don't need to fix everything. In fact, you run the risk of making things worse if you run around trying to fix processes, procedures, and systems. Yet, you need and want to fix the systems that are setting workers up to fail at your facilities.

Here are some quick guidelines for identifying potential areas for systems improvement in your organizations:

Look for high consequence activities

These are the tasks you know will cause consequences if these processes fail. You know these tasks involve many high risk, high consequence factors as a normal part of doing the work. It is fair to say that these systems normally work well, and either rarely fail or have never failed. These potential failures can, and will, cause unacceptable consequences for your organization, your people, or the world at large.

These specific tasks either have high risk hazards or high value materials. These systems often have remarkably robust safety programs associated with them. If you are not assessing these programs for the presence of multiple layers of defenses and operational surety, you should start today.

You probably aren't looking at these systems and asking, "When this process fails, what safety defenses will reduce and control the consequences of this failure?" These systems don't need a stronger industrial safety program. You are probably sure that they are quite reliable, as the process exists now. However, what these processes demand is an aggressive systems safety program.

Assume the failure will happen. Then when the failure does happen, you will be prepared for the failure and its consequences. If the program does not fail, you will still be ready for failure to happen.

Look for small signals that can indicate system weaknesses or problems within the normal work processes

Errors, near misses, good catches, close calls—any of these factors could indicate there is a problem, without the actual consequences of the failure.

Safety professionals look on indicators of this type as "gifts." There are many examples of these small signal events happening every day in your organization.

You must monitor low level events in order to understand where your systems are confusing, conflicting, or potentially flawed. Don't monitor these small events because they are mathematically predictive (the more these events happen, the higher likelihood that we will have a serious event); these small events aren't truly predictive.

What these small events do is to allow an organization to "pulse" its processes and systems to identify potential larger system weaknesses or failures. The premise is that these low signal indicators are in reality pointers towards larger system weaknesses and failures.

Look for error-provoking system steps and processes

These types of conflicts are all over your organization. You know these situations: the cases where if the worker were to follow the process the worker would certainly fail. These are the conflicts that arise in any organization's systems that place workers in positions of uncertainty—while the organization assumes that there is clarity of decision. The classic example used for this discussion is about the difficulty of using a purchasing system.

For example, you need to buy a screw for a project you are doing in your job. The screw will cost about 50 cents at a hardware store. To use the company's purchasing system will take three signatures, filling out a couple of forms, and about two hours of time. What is going to make you use the system when you drive home by a leading lumber seller every day? It is actually much easier, faster, and less painful to buy the screw with personal money than to use the organization's systems, and the organization's money.

Look for error likely conditions

Error conditions are everywhere. Look for the space between your workers and the work they are doing for conditions that could cause failure. Identify places where failure could happen because of the worker simply performing work. They are everywhere. Better yet, look for well-designed systems and conditions that direct workers away from error and toward success. You will not be surprised to discover that many of your systems actually do set workers up for failure. Many of our systems are set up to be effective accounting systems, and not good work management systems. These systems are rife with potential failures.

Listen to your workers

Ask your workers where the next accident will happen; you will be surprised by what you will learn. They are brilliant at this CSI task. Your workers know where your system makes sense, works well, and is efficient. They also, by definition and practical experience, will know exactly where the system is setting workers

up to fail. Remember that when you ask workers for this type of input, you must then do something with this information.

Knowing what the workers believe to be failure-prone becomes a contract with the workers to improve this situation. The key concept here is that mitigation of this type of problem is often better and faster than long-term, expensive solutions. It investigates potential failures before they could take place.

Caution is noteworthy here. If you ask workers to tell you about potential failures, you must be prepared to hear honestly, accept, and act on this information, and not justify and protect systems over people. Don't defend the process over the opinion of the workers. They are your frontline experts.

It is no secret: if you want to know how work actually gets done, ask a worker what is truly going on there. They have the day-to-day practical experience.

Ask yourself (and your peers) what keeps you up at night

Like listening to the workers, you are now creating another path towards capturing another perspective of your work environment. This question is as strong as any other tool used to predict potential failure. Have your fellow managers and supervisors tell you what they think is the most risky activity (or part of an activity) on the shop floor.

Remember that risk perceptions change the closer you get to the actual work. In most cases, risk perception decreases the closer you get to the actual work. So, by definition, management's perception of risk is going to be different to the workers' definition of risk. Managers may be away from the work, but what they don't have at the worker level they most certainly make up for at the systems/relationship level. Your fellow managers know how work fits together: the bigger picture. These people can talk about pressure, resources, and worker personalities.

Get real information

Small things that go wrong are most often warning signs of growing trouble deeper inside the organization's systems and processes. This type of data helps to provide remarkable insight into the health of the entire organization; almost everything that goes wrong is some type of indicator, some type of data. Look deeper into the context of what small failure can mean. Failure often has subtle beginnings and plenty of retrospective indicators.

Near misses and close calls are a function of how much your workers trust you and your organization. Don't underestimate how many close calls your organization may be experiencing every day you are operating. The bigger question here is how many close calls and near misses do you know about, and how do you handle the low level events that actually get to your radar screen?

Where you are in the organization assuredly affects the type and the amount of information that you get from the field. For example, a senior manager by definition will have the view from the top—even if the senior manager believes he or she understands the workers' perspective. Just by entering a workspace as a safety professional, you change that workspace.

The goal is to identify and use multiple ways to get honest information from workers. Once you feel you are getting honest feedback, accept it. Process the information that you get—not the information that you wish you were getting—and don't be defensive. It is easy to try and explain what is going on when you receive a worker comment, but don't do it. Accept the honest information deliberately and thankfully. Remember, the worker is giving you a chance to see the work from their perspective. Actually, that is almost a sacred moment and should be considered significant and vital.

Getting real information has more to do with trust and relationship building with those closest to the potential failure than it has to do with record keeping, computer systems, and accounting. Worry how you will manage the information after you figure out how to build trust levels to a place where your workers can speak honestly about systems and processes. Then take this data and act on it.

Interview your workers

Ask workers what works and what does not work within your organization, and then listen carefully to how they answer these questions. Where do your workers believe they are the safest? Where do your workers believe they are least safe? How do your workers create safety while they do their work? What hazards are they skilled at identifying and mitigating, and which hazards are they surprised by? When do your workers have to adapt and improvise in order to do their work? This is a gold mine of data.

Just as you would if an event had taken place, you would bring in small groups of workers and ask them to tell you how work actually happens. You know how work is supposed to happen, how the system was designed to take place, but do you know how the work actually happens? This interview is to help you capture the story of the work that is being done. In capturing the story you will learn of the parts of the work that are safe, and the parts of the work that need your attention.

Mostly, you are interviewing workers to ask them to tell you the story of working in your organization. This dialog will tell you where failures can happen before they happen.

Investigate the accident that hasn't happened…yet

What you go out and look for is what you will go out and find. Be careful. You must be acutely aware of how your pre-investigatory biases will directly impact what you will find. Since there has been no event, the normal "smoking ashes"

that you would sift through are not present in a pre-accident investigation. This forces you to be imaginative and creative in the way you proceed in this learning opportunity. You are honing your CSI skills.

In short, what you want is not an investigation, but in reality an explanation of how failure would happen, if it were to happen, in the work being done by workers. You are writing a story of what could happen if the organization did not intervene in the process. A story is a powerful tool because it has a beginning, middle, and an end. A story must make sense, be plausible and realistic, and move logically through time and space. You can sensibly explain what could happen, and then how to prevent it.

This is the time you will capture information at a detailed level. The first thing you will want to do is to start building a timeline. The difference between this timeline and an actual failure timeline is that your timeline for a pre-accident investigation does not have an accident on the right end of the chart. What this timeline has is an intervention and defense placement on the outcome end of your chart.

Never look for cause and effect outcomes, instead look for relationships between workers, workers and management, and, most importantly, workers and your systems and technologies. Looking for these relationships forces your organization to understand and "deal with" the context of your work. It is within this context that you discover the potential failures before they become failures.

Now, engage your learning to strengthen your systems against the potential failures you have uncovered

You're not ever going to be able to stop an accident. You can directly change the way the accident affects your organization, your workers, and yourself. A pre-accident investigation helps you make your organization better prepared for a failure. Our message is consistent and straightforward: "We can't stop accidents, but we can prevent consequences."

Engage your learning systems to make your organization smarter and more prepared for the potential failures you uncover.

You are gathering information to prevent adverse outcomes for your workers and your organization. You will never be able to measure what doesn't happen. You will never be able to predict every event. Yet, it is clear that if you can gather enough information about a system to identify the places where failure is most likely to happen, or places where if a failure were to happen it would have some type of serious consequence, you can actually intervene in your organization's processes and systems, and prevent events.

You aren't actually looking for the next accident. You can't ever predict the next accident, because these events are always surprises to the worker and to the organization. What your job becomes is to identify places within your organization where the organization sets up workers to fail. Look for conditions

that lead to failure, not trends or numeric patterns, and when you find these conditions, learn from them.

It all boils down to how your organization, your managers, your teams, and you learn from your organization's normal, routine work—day in and day out. The only tool you have to prevent events from happening is your organization's ability to learn as an organization—to learn honestly from itself. Use your learning systems to serve the greater good, not to function as the accounting method to document your safety history.

You will struggle to look for indicators that you can use as predictive. Every organization struggles to identify indicators that can be used to predict the next critical event. I am convinced that there is no such thing as a *leading* data set that can help safety professionals predict the future. I can make this claim because if this data set were out there, I am positive that someone would have found it, and we all would be using these numbers and methods.

Most importantly, ask people what conditions we should identify. Where are we most prone to having a problem? Where are we sending conflicting, multiple signals? Where is work overly controlled? Where is work under-processed, under-resourced, or misunderstood? Ultimately, use your information to the best of your ability.

How will you know you are correct? How will you know you have been effective? The quick answer is that you never will know how many events didn't happen because of your pre-accident investigation process. You will never be able to measure something that does not happen. Be satisfied in the process, not in the outcome. After all, if you have done your job well, nothing unexpected will happen. What a beautiful thought for your workplace—imagine, a workplace where nothing unexpected happens every day.

Case Study
Aviation Accidents are the Unexpected Combination of Normal Aviation Variability

"When everyone is to blame, is any one person blameworthy? Does blame have any value at all in better understanding this event?"

This case study demonstrates how almost no indicators of failure are present before the deviation from the expected outcome transpires. Perhaps even more interesting is the rather normal nature of all of the factors involved. No one factor seems important alone. No one indicator is even considered operationally interesting by itself. Had there been any indication, at any level, that this outcome was even possible, many different decisions would have transpired on the night of December 20, 2008. It typifies a difficult conundrum of safety management: no predictable indicators of failure *before* consequence happens.

While reading this case study, observe how this event is made to appear linear (it is actually the only way to talk about this event in retrospect). Contrast this same linear failure with the actual failure that was the outcome of multiple conditions all coinciding nearly at the same time. You will see how these conditions that are cited in this event all seem to happen repeatedly within the telling of this story. It is important to note that this case study is the product of an aviation insider telling a context-rich story of a successful accident.

Finally, notice how the absence of any formal operational leadership (due to injuries sustained by the captain and first officer) after the event gave rise to the opportunity for leadership to move throughout the event's context. Leadership for event recovery, event response, and passenger evacuation moves many times throughout this event.

THE BURNING OF FLIGHT 1404

On Saturday evening, December 20, 2008, flight #1404 crashed on take-off from a runway at Denver International Airport, bound for George Bush Intercontinental Airport, Houston, Texas. The accident resulted in many

passenger and crew injuries, and a hull loss of a Boeing 737-524 aircraft. According to the preliminary information provided, the plane veered off the runway early in the take-off roll, ultimately crashing into a ravine approximately 200 yards from the runway center line and erupting in fire. The aircraft sustained severe damage. The fuselage was cracked just behind the wings. Engine number one and the main landing gear were sheared off. In addition, the nose gear collapsed. The fire caused overhead luggage compartments to melt onto the seats. The fire was on the right side of the aircraft. All on board evacuated through the doors on the left side of the aircraft.

The initial emergency dispatch call to the Denver Fire Department was incorrect in the location of the aircraft. However, despite the early confusion as to the whereabouts of flight 1404, firefighters were on the scene relatively quickly, as the aircraft came to rest near one of the airport's four firehouses. When the Fire Department arrived, most of the right side of the plane was on fire. The passengers were evacuating from the left side, assisted by flight attendants and one off-duty airline pilot in the passenger compartment. This same pilot made several trips in and out of the wreckage to ensure everyone was safely out of the aircraft.

Initial evaluation of the accident cited rudder issues as the reason for the aborted take-off. Further investigation indicated issues involving a combination of high crosswinds, or wind shear, and the presence of aftermarket winglets installed on the wings. An additional reason was the captain's inability to keep the aircraft on the runway during take-off. A plane that leaves the runway during take-off is said to have had an runway excursion. Try to keep all of these reasons in mind as you sift your way through this case study.

This was the most serious accident in the thirteen-year history of Denver International Airport. Luckily, there were no fatalities. However, there were some notable injuries.

A SUCCESSFUL ACCIDENT?

Flight 1404 was scheduled to depart Denver for Houston at 18:00. The flight crew arrived at the airport about 17:00. The airplane had not yet arrived at the gate. The first officer bought coffee while the captain walked downstairs to get flight paperwork from an operations coordinator.

The flight crew consisted of two pilots: a captain and a first officer. The day of the accident was the fourth day of a four-day pairing for the two pilots. The pilots had been paired with each other on one or two prior trips. The last trip had been about a month before the accident. The pilots reported a history of positive, professional interactions. It was a typical day and a typical flight.

The pilots met at the gate after the airplane arrived. The captain did an external preflight inspection, while the first officer performed preflight safety

checks in the cockpit. When the captain returned to the cockpit, the two pilots discussed the upcoming flight; performed the receiving aircraft checklist; obtained a passenger count, and weight and balance information; and entered the load information into the airplane's flight management computer. All these are standard preflight tasks.

The flight crew instructed the cabin crew to close the airplane's doors. The first officer contacted ramp control and received a clearance to push back from the gate for a west taxi. The flight pushed back at 18:01. Ice and snow were visible on the ramp. Following standard practice, the captain started both engines, and turned on the engine and wing anti-ice systems. The first officer contacted ground control, and received a clearance to taxi to runway for their departure.

The flight crew heard ground control tell the flight in front of them that Automated Traffic Information System (ATIS) "Sierra" was current. Winds were reported at 270 degrees and 11 knots (about 13 mph). They continued to taxi on the taxiway toward the runway. The captain did not notice any buffeting of the airplane from wind during the taxi.

As the airplane approached the runway, the flight crew performed the before take-off checklist and contacted the tower. An additional aircraft was on the runway ahead of them, awaiting a take-off clearance. When the awaiting plane departed, the tower instructed flight 1404 to position and hold on the runway. The runway appeared to be clear of snow and ice, so the captain decided to deselect the engine and wing anti-ice systems, but he left the engine igniters on. The captain positioned the airplane on the runway, and the flight crew waited for two or three minutes. The cabin crew waited in their seats. The passengers were thumbing through magazines, adjusting pillows, and talking amongst themselves.

The runway lights and all of the airplane's lights were operating. Visibility on the runway was excellent. The captain later recalled that the airplane's speed was probably less than 100 knots (115 mph) when this event occurred, as he had not yet heard the first officer's 100 knot callout.

The captain received no system warnings and no signs of obvious system malfunctions. He kept the right rudder fully engaged as the airplane veered toward the edge of the runway. His right hand remained on the throttles and his foot continued to depress the right rudder pedal during this time. He did not touch the brakes because he did not want to interfere with the auto-brakes, which were selected to the "rejected take-off" setting.

The tower contacted the flight crew, informed them that winds were 270 degrees at 27 knots (31 mph), and cleared them for take-off. The controller's wind report surprised the flight crew because it was higher than the 11 knot winds reported in ATIS Sierra. The captain recalls telling the first officer something like, "Roger, crosswind." The first officer recalls the captain saying, "Winds are 270 at 27. Are you ready?"

The captain began a reduced power take-off. He first pushed the thrust levers up and then increased power. He noticed a difference in the thrust being generated by the two engines, but the two engines matched as he increased to 90 percent. After verifying this, he pressed the TOGA (flight deck intercom) button and called out, "check power." The first officer responded that thrust was set at 90.9 percent. The captain applied a left control wheel correction, applied forward pressure to the yoke, and used variable right rudder to keep the airplane aligned with the runway center line. It felt at first like a normal crosswind take-off.

As the airplane was getting up to speed it suddenly swerved to the left, as if hit by a gust of wind, or as if the tires had hit a patch of ice and lost traction. The captain used full right rudder, but saw the airplane continue to veer left. As airspeed was increasing from 87 to 90 knots (103 mph), the first officer looked up and saw the airplane drifting left of the runway center line. He thought the captain was correcting back to the right, but the airplane suddenly yawed 30 to 45 degrees to the left. It appeared to the first officer that there was no directional control. He felt the rudder pedals with his feet and he believed the captain was applying full right rudder.

The captain saw the edge lights on the left side of the runway. He believed the airplane was going to exit the left side of the runway and, as a last resort, he reached down with his left hand and grabbed the tiller for a second or two. He attempted to steer the airplane back onto the runway using the tiller, but this did not work so he put his left hand back on the yoke. These decisions were quick. It was one of those moments in life when time feels as if it is moving in slow motion, but in actuality is moving lightning fast.

The captain used the right control wheel to keep the wings level as the airplane departed the left side of the runway. As he did this, he thought the ground next to the runway sloped down. He feared the aft end of the fuselage would slide down that incline and cause the airplane to tumble on its side. After the airplane had thoroughly exited the runway, the captain said, "reject" and tried to deploy the thrust reversers. He was unable to deploy the reversers because the ride was incredibly rough. The plane was shaking. Cabin lights were flickering. The passengers were jolted and surprised by the drama transpiring.

The airplane, crew, and passengers suffered two violent impacts before the plane came to a stop. It was totally dark in the cockpit. Neither the captain nor the first officer heard any engine sounds. Both were stunned and injured. They felt incapable of doing anything for what seemed like one or two minutes. They did not order the standard evacuation, nor did they begin the evacuation checklist. They had trained in these procedures and drilled them many times. They were unable to perform them at that moment. The passengers were momentarily leaderless.

Both pilots recall that the flight deck was pitch-dark. The first officer said that he could hear things going on in the cabin. He thought that he needed to make a public announcement, but did not. His next thoughts were about getting both himself and the captain out of the aircraft. As soon as he recovered from the initial shock of the crash, the first officer opened a cockpit window to his right and threw out an escape rope. He saw fire along the right side of the airplane. He got out of his seat and exited through the cabin instead.

Evacuation

Meanwhile, back in the passenger cabin, the damage inside the aircraft was considerable. A fire was quickly consuming the right side of the aircraft, worsening the situation. Upon looking out the windows and identifying imminent danger, the flight attendants commenced an evacuation, and ushered the passengers out of the burning plane with the help of two deadheading pilots. ("Deadheading" is a term in the aviation industry which refers to airline flight staff who are being transported free of charge and are not working on the flight.) The deadheading pilots did not go out the right side of the aircraft because of the fire. The cockpit door remained closed for the entire evacuation. By the time the first officer opened the door, everyone had exited the aircraft except the deadheading crew and one of the flight attendants.

As the first officer stood up, a deadheading crew member knocked on the cockpit door. The first officer opened it. About this time, the captain was trying to get out of his seat as well, but a dislodged flight crew bag was blocking his path. The first officer moved the bag, and with the assistance of the deadheading crew member helped the captain out of the cockpit. The deadheading crew member told the first officer that all of the passengers had been evacuated, and then the captain, first officer, and deadheading crew member exited the airplane via the door slide.

The three flight attendants experienced no problems with the escape slides, and the emergency exit lights were brightly illuminated, which proved quite helpful in guiding the passengers out of the aircraft, given the dark conditions. All passengers and crew exited the airplane via the left side doors and overwing exit. The flight attendants reported that the passenger who opened the overwing exit did so quickly and easily, in part due to the detailed safety briefing given to each passenger in the emergency exit row seats prior to every take-off. A bottleneck of people developed by the left overwing exit, and the deadheading captain directed passengers out via the doors or floor exits.

The cabin was illuminated because of the fire. The entire wing and wing root were on fire, which was most noticeable at the overwing exit. The deadhead first officer unbuckled his seat belt, turned left, and saw that

the male passenger sitting in the exit row had got the door open quickly. There was a tremendous confluence of passengers trying to exit through the overwing exit. Five people were trying to get out first; no one wanted to be second. The windows were melting and popping. Passengers were screaming, "we're gonna burn" and "it's gonna explode." It was incredibly frightening for the passengers. Many people were trying to get out at the same time; some passengers were climbing over the seats to exit. There was so much panic that most instructions fell on deaf ears.

The other deadheading crewmember, a first officer, assisted several elderly ladies in evacuating, and re-entered the aircraft. He noticed the two pilots emerging from the flight deck being helped by the deadheading pilot, both obviously in pain. He helped them to exit through the door, and then came back for the first class flight attendant to help her, since she had an injured ankle.

The deadheading first officer assisted the injured flight attendant off the plane, and returned into the plane. He found the deadheading captain and the male aft galley flight attendant in the aisle as the plane started to fill with smoke. They met in the middle, over the wing, and began to check for anyone else who might have been left on the airplane. There had been a large number of small children on the flight. The two wanted to make sure everyone had evacuated. The aft flight attendant said it was all clear in the back. The deadheading captain made sure by asking if everyone had departed; the deadheading first officer said yes, and the captain told him to go back and check one more time.

By this time, fire and smoke were starting to come up through the floor. The crew were concerned the center fuel tank might explode. By the time they left the aircraft, the windows were starting to melt, and they were fearful there would soon be a breach; everyone knew it was time to get out of the aircraft.

Outside the aircraft, the first class flight attendant was in a great deal of pain and could not stand up. The deadheading first officer picked her up and carried her away, as the fire was growing fast. The center tank gave way, and a river of fuel ran north–south toward the nose, and fire was following behind it. It was starting to light up the entire cabin.

Luckily, the slides deployed properly. Since the landing gear was sheared off, the doorsills were not far above ground level. The evacuation slide was more a padded walkway. However, the gray colored slides might have been easier to see if they had been bright yellow or fluorescent orange.

Runway Conditions at the time of the Event

Inspection of the runway following the accident revealed that it was bare, dry, and free of debris. The first tire marks were found about 1,900 feet from the runway threshold. The aircraft exited the runway about 2,650 feet from the runway threshold, and continued across a snow covered grassy drainage basin area. It then crossed a taxiway and a service road before coming to rest about 2,300 feet from the point at which it departed the runway. On the night of the accident, the majority of air traffic was arriving at the airport from the south, and departing from the airport to the north.

Weather Conditions at the time of the Event

Both pilots were aware of the crosswind conditions, having been advised by Air Traffic Control that winds were 270 degrees at 27 knots (31 mph) before take-off. The weather observation in effect for Denver International Airport nearest the time of the accident was reported to be winds at 290 degrees and 24 knots (27 mph), with gusts to 32 knots (36 mph), visibility of 10 miles, a few clouds at 4,000 feet, and scattered clouds at 10,000 feet. The temperature was reported as 24 degrees Fahrenheit. Wind data had been obtained from the airport's low level wind shear alert system, consisting of 32 sensors located around the field which recorded wind speed and direction every 10 seconds. This information would have been used to determine a better estimate of the actual crosswind component at the time of the accident.

Leadership

The captain was interviewed four days after the accident while still hospitalized with injuries that included spinal fractures. He stated that he "was either knocked out or dazed" immediately after the crash, and did not recall how he got out of the airplane. Both pilots recalled that the flight deck was totally dark. The first officer said that he "could hear things going on in the cabin," and he thought that he "needed to make a public announcement" but he did not. His next thoughts, he said, were about getting himself, and the captain, out of the aircraft. The cockpit door was closed for the entire evacuation. By the time the deadheading first officer opened the door, everyone had vacated the aircraft except the deadheading crew and a flight attendant. Meanwhile, back in the passenger cabin, the damage inside the aircraft was considerable, and a fire that was quickly consuming the right side of the aircraft was worsening the situation.

As soon as the aircraft came to a stop, the deadheading captain unbuckled his seat belt, and, although injured, assisted with the evacuation that was already underway. He gave a description of the conditions inside the cabin. He said the panels in the middle of the row "had swung down and were still swinging." He tried to keep them out of the way as people went by, because he knew "the passengers could be injured because they swung so fast." Although passengers were bumping him as they passed by in the aisle, he finally got the

panels up and locked into place. He then went toward the back and got three more panels locked up.

He said they hinged on the aircraft right in the aisle, and they were down and swinging back and forth. He pushed them back in place, and that was how he was hit. He said he was 6' 3", and was holding them back. Passengers holding babies hit him a couple of times. He jumped to the other side, and pushed a panel up, got hit by another passenger, and fortunately locked the panel back into place. He said the panels did not malfunction, they had just come undone.

As the fire grew more intense, the deadheading captain could see a breach in the cabin just aft of the exit row. He saw the emergency lights were on, but could not see past the breach because it was dark in the back. He saw the flames on the outside of the aircraft through the windows to the overwing exit. He said he did not feel any heat initially. The deadheading captain then looked toward the cockpit. He saw "the forward flight attendant was standing on one leg, holding herself up. The cockpit door remained closed." The forward flight attendant assisted with the evacuation of all the passengers. Then the deadheading first officer opened the cockpit door. The captain "was out of his seat between the pedestal and the cockpit door, and was in excruciating pain." The deadheading captain said he could see that "both pilots were seriously injured," and that "they looked dazed from the impact." The two deadheading pilots got the flight deck crew out of the cockpit and helped them through the door. The fire was on the right side of the aircraft. All on board evacuated through the doors on the left side of the aircraft.

Crew Mindset and Context

The accident flight was the first flight on the fourth day of a four-day trip for the crew—comprised of the captain, the first officer, and three flight attendants. The captain, the pilot flying, had accumulated a total of about 13,000 hours, with about 5,000 in the 737. The first officer had flown about 7,500 hours in his career, with about 1,500 hours in the 737. The first officer was the pilot monitoring. All appeared normal until the aircraft began to deviate from the runway center line. The captain noted that the airplane suddenly diverged to the left, and attempts to correct the deviation with the rudder were unsuccessful. He briefly attempted to return the aircraft to the center line by using the tiller to manipulate the steering of the nose gear, but was unable to keep the aircraft on the runway. Additionally, there were two of this specific airline's pilots who were on board as passengers while deadheading back to their base—the deadheading captain and the deadheading first officer. The two had flown the aircraft into Denver on the previous flight, and were passengers on the accident flight. Neither pilot cited any anomalies on the

inbound flight, and they reported that all the aircraft's systems had operated normally.

Data Recorders

Bumping and rattling sounds audible on the Cockpit Voice Recorder have been time-correlated with the Flight Data Recorder, and were found to have occurred as the airplane exited the runway and traveled through the grassy areas adjacent to the runway. The aircraft reached a maximum speed of 119 knots (137 mph). It was traveling at 89 knots (102 mph) when the Cockpit Voice Recorder and Flight Data Recorder stopped recording. Physical inspection of the engines and information from the Flight Data Recorder did not indicate any evidence of pre-impact malfunction in either engine. The Flight Data Recorder data indicated that the number one engine power was reduced before that of the number two engine during the accident sequence, and examination of the engine indicated this reduction to be consistent with snow and earth ingestion as the airplane departed the runway. The Flight Data Recorder data also showed that both engines had been commanded into reverse thrust following rejection of the take-off by the flight crew, which occurred after the aircraft had already left the runway.

Landing Gear and Controls

A preliminary examination of the rudder system revealed no abnormalities or malfunctions. The main landing gear and brakes, which had separated from the aircraft during the accident, were found to be in working condition, as there were no signs of the hydraulic systems leaking or flat spots on the tires. The flight deck controls and corresponding control surfaces were found to be in the take-off configuration.

Aircraft Cabin

There was no obvious damage to the passenger seats, which were found secured to their tracks. The safety belts were all intact, although some showed evidence of fire damage.

Injuries

Of the 110 passengers and the 5 crew on board, some 38 sustained injuries, including broken bones, though everyone on board survived. Two injuries were critical, though both passengers' medical conditions were upgraded that evening. By the following morning, fewer than seven people remained hospitalized.

The captain was one of the critically injured. He was hospitalized with serious back injuries and bone fractures.

Investigation

Despite conflicting initial reports, and reported media inaccuracies, this accident is considered a "successful accident" due to several contributing factors: the flight crew's immediate actions, the presence of the deadheading airline employees on their way home, the fire rescue training that was developed and practiced for just such an event, and divine intervention or luck.

Incident response timeline:

1. Crash event happens at 18:17.
2. Crash reported at 18:18 local time by the air traffic control tower.
3. Initial dispatch—inaccurate location and limited information.
4. Denver Fire Department was on the scene of the actual accident at 18:23 to extinguish the fire.
5. Denver Fire Department confirms exterior fire out at 18:25; moves to fighting interior fire.
6. Denver Fire Department reports all fire out at 18:28 and initial primary search results are negative.

The aircraft's black boxes were recovered from the wreckage in usable condition. The Cockpit Voice Recorder did not reveal any apparent problem until 41 seconds after the aircraft's brakes were released, just before take-off. At that point, a bumping or rattling sound was heard, and the crew aborted the take-off four seconds later. Both recorders stopped working 6 seconds after that (before the plane came to a stop). At one point during the sequence, the plane's speed reached 119 knots (137 mph). Both the captain and first officer had outstanding safety records when the crash occurred, and were experienced pilots. Wheel marks left on the ground as well as initial reports from passengers and firefighters indicate that the plane was airborne, briefly. It is unclear at which point during the sequence the fire started. There was no snow or ice on the runway; however, there were 31 knot (36 mph) crosswinds at the time.

The flight crew that flew the aircraft to Denver before the incident flight were also on board, though not on duty, and reported having no difficulties with the plane during their previous flight. The aircraft was clearly in operational order before the runway excursion event happened.

Initial reports indicated that the plane could have suffered a landing gear malfunction that might have resulted in a wheel lock-up during the take-off roll, leading to the runway excursion. National Transportation Safety Board officials said that when the take-off began, the aircraft's engines appeared to be functioning properly, its tires were inflated, and the brakes did not appear as if they had failed or otherwise malfunctioned. They concluded that the landing gear had not caused any problems.

On July 17, 2009, it was announced that focus had shifted to a possible large gust of wind or a patch of ice. The pilot of the aircraft stated that: "My speculation is that we either got a big, nasty gust of wind or that, with the controls we had in, we hit some ice." The winds were reported at about 24 knots (28 mph) to 26 knots (31 mph) from the north-west, with gusts up to nearly 32 knots (37 mph), just before the airliner began its take-off roll northward down a north–south runway. The 737 is specified by design for a crosswind limitation during take-off of 33 knots (38 mph) on a dry runway.

According to an article published on July 18, 2009, in the *Denver Post*, flight 1404 had been equipped with "winglets"—curved, upswept structures added to the tips of wings—in November 2008. The article states: "The B-737 flight manual 'limitations section' showed a crosswind limit of 33 knots for a dry runway. But NTSB added that the manufacturer and installer of winglets that were on the airplane had published a maximum demonstrated crosswind component of 22 knots (35 mph) for winglet-equipped B-737-500s."

On July 13, 2010, the National Transportation Safety Board published that the probable cause of this accident was the captain's cessation of right rudder input, which was needed to maintain directional control of the airplane, about 4 seconds before the excursion, when the airplane encountered a strong and gusty crosswind that exceeded the captain's training and experience. Contributing to the accident was an air traffic control system that did not require or facilitate the dissemination of key available wind information to air traffic controllers and pilots, and inadequate crosswind training in the airline industry. This example illustrates surprisingly basic deficiencies in individuals' skills and in organizational capacity that pilots and air traffic controllers were unaware of as recently as 2008.

Interestingly, had this rejected landing event happened 100 yards further down the runway, the event would have happened directly in front of one of the airport's firehouses. Even though the event did happen close to a firehouse, the fire rescue teams were delayed by mistakenly being told when dispatched to head out of the firehouse and turn to the left. Had the fire rescue teams turned right instead of left, they would have been at the event nearly immediately.

When the fire rescue teams rolled up on the scene of this runway reject event, the firefighters doing the initial assessment of the accident noticed the pilot in the flight deck of the plane waving his arms towards the responding rescue crew members. The firefighters assumed that the waving arms were a signal that the pilot and first officer were OK—there were no problems in the flight deck of the plane. In fact, the pilot was attempting to signal to the rescue

team that they needed assistance, and that there was an identified problem at the front of the plane. Is there a difference between waving for and waving off assistance during an event?

DISCUSSION

Let us identify some of the conditions that, although normal, combined to create this DFEO:

- Weather
- Skills of crew
- Warning and notification systems
- Quality of operational information
- Surprising outcomes happening simultaneously
- Design features of the wings—which create clear fuel efficiency
- Emergency response systems
- Operational leadership
- Normal system responses up to the point of the failure

This event is "classic" in the unexpected combination of relatively normal aviation performance variability, which produced a clear, and potentially life-threatening, deviation from an expected outcome—in this case a safe and on-time departure from the Denver International Airport. All of the conditions listed are extremely beneficial in creating the context for this failure. Keep in mind that this by no means is a complete list. The problem is that all of these conditions are also incredibly normal.

4
Workers Don't Cause Failure, Workers Trigger Failure

"Cause is not found in the rubble. Cause is created in the minds of the investigators..."

Sidney Dekker, *The Field Guide to Understanding Human Error*

QUESTIONS ABOUT YOUR ORGANIZATION:

- How long would it take to find a worker at my organization not following a procedure?
- Does not following an organizational procedure always create an immediate failure?
- Why is it when we have an event that we attribute that event to a failure by the worker to follow procedures?

REMEMBER: THREE PARTS OF A FAILURE

Every event you will have or have had in your organization is actually made up of three distinctly different parts:

1. **The Context of the Failure**: This is everything that led up to the actual event. This is the worker, the worker's mindset, the environment, the people, the hazards, the planning, and everything else that make up the story of the failure. Remember, in the context section of the event there is much information that the worker does not or cannot know—for example, how the event will end.
2. **The Consequence of the Failure**: This is what happened, the effect of the failure, the damage or harm, the deviation from the expected behavior. This is the ending of the event.
3. **The Retrospective Understanding of the Failure**: This is how the organization views the failure or deviation from the expected behavior. It is at this point that all knowledge is available about this failure. The retrospective has almost superhuman knowledge of every facet of this event. This is when most organizations look for a cause—a specific

moment in time when something bad happened—and use that cause to understand why the failure transpired.

CAUSE IS A CONSTRUCT—CAUSE IS MADE UP

How crazy is this going to be? I am about to commit safety blasphemy. I am going to tell you that workers don't cause failure. It is true, workers don't cause failure; what workers do is trigger a whole lot of weaknesses that exist in environments, processes, systems, job sites, and in the work or organization itself. These weaknesses are not new or unusual. These weaknesses are always present in your organization.

That is right—workers don't cause events. Seriously. But the problem with this statement is not what you think—it is not the reduction of accountability for failure from the workers. We are not trying to create an organization that gives all workers a free pass. Workers are always accountable for their actions—every single action. Just as companies and organizations are also always accountable for their systems and processes that bind and drive worker actions. The problem is the use of the word "cause."

Here is the problem. Culturally, we have become trapped by the idea that everything must have a cause. Most organizations use the term "root causes" when they do failure investigations. We call them "causal factors" when we thoroughly explore failure. We think if we can just find the cause of the event, and fix the cause of the event, our workplaces will be safer and our problems will all go away.

In the Freon case study at the start of this book, what is the cause of that event?

Could the cause be a stupid worker? Could the cause be a lack of supervision 8 hours a day, 5 days a week? Could the cause be the presence of Propane tanks in the Freon storage area? You read the case study. What was your idea of the cause for that event? Is there just one cause? Are there many causes? Is there a root cause? Is the cause in the Freon case study even necessary in understanding that failure? What improvement actions can we apply, based upon what we understand from this search for cause?

Organizations need cause. They must have something or someone to blame for such a stupid act. It's as if you cannot operate your facilities if you don't know what caused the most recent failure. Why is there such a strong pull to find a cause and fix the problem? Let's explore this idea.

Why do organizations have this almost zealous need to find the cause of a failure? The following reasons, listed in order of perceived operational need, may help answer this question.

You must have a cause for every event that happens at your facility because:

1. **Not knowing what caused an accident is downright scary.** It is simply not acceptable not to know how something happened. You can't tell your management, "we stopped being 'lucky' and, therefore, had a failure." You would go crazy and you would get fired from your job.
2. **Headquarters, management, or the regulator require a cause statement.** The regulators at my facility want to know three things immediately after a failure happens. Who did it? How did we let them do it? And where are we going to transfer that person, so it never happens again? If those questions are what the regulator wants to know, it is no surprise that the answers to those three questions are what we look for when we investigate and write up a failure.
3. **We have a need to punish the guilty.** You could also say that we have a strong need to tell the rest of the world how, had it been me doing that work, "I would not have done the same things the worker did in order to cause the failure." A better worker would have been smarter, or cared more, or not been so stupid. This need to punish the guilty seems to be almost instinctual in human beings. We probably don't want to punish the guilty; actually we want to assure that the failure that just happened cannot happen again. The only way we can do that to our satisfaction is to attribute the failure to the guilty person, and not to luck or to the system. We humans are a complex bunch when it comes to assuring stability in our lives and systems.
4. **We must develop corrective actions.** We must do something as soon as possible to bring our systems back in order. One of the strongest drivers to find the bad worker and identify a root cause is the fact that it is much easier and cheaper to fix *one* worker and fix *one* cause than it is to look at an entire set of processes and relationships that exist throughout an entire organization. Be careful not to fall for the easy answer—it has the potential to create a less reliable and more dangerous system, not performance success.
5. **We want to know how to adjust our organization's systems and behaviors to avoid the same kind of trouble happening again.** Finally, here is a reason I can stand behind. This is exactly how you should always analyze every type and every level of failure. Ultimately, you should only try to understand failure for what you can learn, not for who you can blame. It's too bad that this is the fifth entry on our list and not the first. Too bad that we must fight past the "eureka moment" of the human failure in order to get to the deeper issues, and discover the fundamental system weaknesses and conflicts. In many ways we construct cause in order to make our need for the creation of improvement actions quick, simple, and submitable.

Therein lies your challenge. All of the factors listed above for the Freon event (plus many more) caused the event to happen, and, in a way, all of those listed factors did *not* cause the Freon event to happen. Failure is the sum total of all of the things that exist in a specific worksite, and all of their relationships with each other. None of the individual parts of the Freon failure were sufficient to cause the overall failure—but all of these individual parts were collectively sufficient to cause the failure.

In a classic root cause analysis (RCA), your job is to deconstruct the event down to its minutest parts, analyze those parts, and fix whatever is broken. In Human Performance, you do almost the opposite. Instead of deconstructing the event, you construct the event context, and look not at the individual pieces of the event but at the relationships between those pieces.

I am getting a bit ahead of myself.

In order to build my case against cause, I have to have a small discussion around the way we solve problems. Before you jump to discredit this discussion, please allow a little conversation about how "cause" forces an organization down one specific path (mechanistic), while looking beyond cause allows our organizations to go down another, different path (holistic).

You must also talk about and introduce the concept of human error. We have discussed error all the way through this book, and we have discussed it on purpose. In order to have even half a chance of moving your organization away from the classic cause and effect trap that we all find ourselves in, we must give managers and leaders in the organization a logical reason to question cause. Knowing that errors happen even when a worker does not mean for an error to happen allows you to have a discussion about why human error is neither a cause nor an effect. More importantly, there is no relationship between error and effects. Error just happens; it either does or does not have an effect. Cause, according to the rules of logic, always creates effect. Error is not a cause, nor is it an effect. Error is simply error—an unintentional deviation from an expected behavior or outcome (DFEO).

MECHANICAL FAILURE VERSUS HUMAN FAILURE

Causal analysis is the best tool in the entire world to find and understand mechanical failure. There simply is nothing better. Cause also helps describe the laws of physics incredibly well. So let's begin this argument by saying straight out, "if you have a machine that breaks down, and you want to know both why it broke and how to fix it, use all the classic causal analysis tools in your toolbox." Those causal analysis tools will work perfectly, and you will get the exact result you want and need.

If my car breaks down, I want my mechanic to use causal analysis to determine what is wrong with the car. Even better, I want my mechanic to identify the "root cause" of the car's failure, then either fix or replace the part in question, and get my car up and running, in "factory floor" condition.

However, if your organization has a failure that includes people and technology, or people and automation, or people and machines, causal analysis tools are not the best choice in understanding "how" the failure happened. As a matter of fact, for those events involving relationships between people and machines, or people and technology, causal analysis is the wrong tool.

I am sure I don't want my doctor to use the classic root cause analysis procedure to fix me if I am broken somehow (nor would you want your doctor to do the same with you if you were broken). If I get sick, I want a doctor who thinks of everything that could be going on with my health, then slowly and surely examines me as a complete system. I certainly don't want my doctor to take me apart and look for what is wrong with me. If that were to happen, I would be more concerned with the doctor putting me back together.

Instead, the doctor tries to construct their diagnosis based upon a series of small (or sometimes large) symptoms that I may be displaying. Some of the symptoms I will tell the doctor I have, other symptoms the doctor will discover during closer examination and tests. All of these individual symptoms will now be analyzed holistically to lead the doctor to a complete way to diagnose what needs to be done to make me better.

The safety world has surrendered a lot of its ability to understand failure to a classic, cause and effect understanding of the world. That has had negative consequences for our organizations and for our job sites. It is a terrible thing that we don't look for and construct the whole story. Instead, we search until we find the first thing broken—often called the "eureka component." We call that first broken thing out and fix it.

Here's the challenge: that classic world view is not complete enough, so we need to make it big enough to capture all of the failure. Cause and effect implies a remarkably linear failure path, a mechanical failure. Human failure is never linear. Human failure always involves complex relationships: relationships between people and processes, people and technology, people and machines, people and other people.

One of the best places to show the importance of looking at the event as a system warning is found in the way we think of jobs for workers in the field. There is a vast difference between the work we think we give workers and the work the workers actually discover in the field.

THE DIFFERENCE BETWEEN WORK AS PLANNED AND WORK AS DISCOVERED

Your organization plans work. You scope contracts. You manage available resources. You have your daily planner, list, and calendars; you plan the crap

out of every possible thing you do. It is a smart way to do business, and a powerful way to ensure mission success and reduce operational surprises. The only problem with having a carefully developed plan happens when the human worker actually starts following the plan. We all know that a plan is perfect until we begin to use the plan.

The following graphic helps guide this discussion. The solid line—the work as imagined or work as planned line—represents work performed the way everyone in the organization believes work should happen. Think of the work as imagined line as perfect work, perfect customers, perfect workers, perfect weather, perfect tools, and perfect processes. Everything that can go well is going to go well, and usually does.

If the solid line represents a perfect world, the dotted line represents reality. In fact, the dotted line seems to be much more representative of how work truly is to the worker. Our workers discover work; the workers almost never are presented with the work they are told they will get. And in order to be successful, workers must learn to adapt your organization's processes and procedures to match the work in front of them. Not vice versa. Workers can't change the work; therefore, they must change the processes by which the work will be accomplished.

The solid line—doing only what has been planned to happen—does not lead to mission success. The dotted line—adaptive behavior exercised in the field—is *always* how work gets accomplished. In fact, the dotted line is the only way workers ever create mission success. If your workers did not adapt, your important work would not get done. Given these ideas, it is vital to emphasize that the dotted line is not a terrible thing. You should not approach the dotted line as something you need to fix, or align more closely with the solid line. The

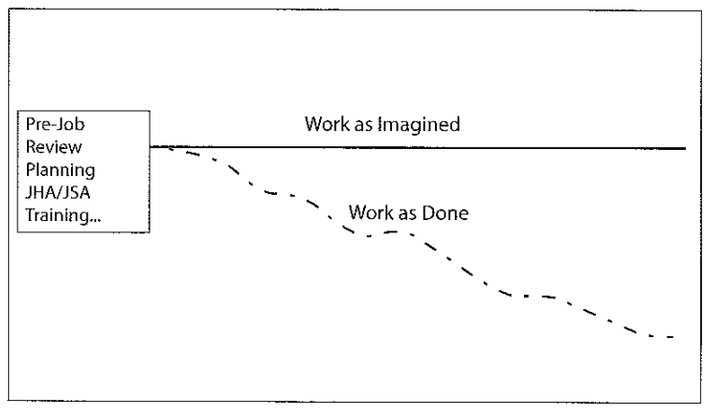

Figure 4.1 Adaptive Nature of Work as Performed

dotted line, work as reality, is always a part of the work. This two-line model precisely represents every job you have ever done. Since we know this model is always present when workers do work, we should use this model as a tool for performance improvement.

The space between planned work and performed work is the operational gap. In that operational gap lives a vast amount of information. This is where you learn about safety, as safety exists in your operations. This is very different from observing worker behavior or auditing procedural use and adherence. This is real post-job information about what happened when work was being done. Finding and understanding the difference between "work as imagined" and "work as actually done" is like finding the place where all your safety data "hangs out." This is a treasure trove of information. Seek understanding here in order to know how to better plan work the next time you perform this task, or tasks like this task.

We *plan* for safety in the pre-job activity. We do hazard identification work, job hazard analysis (JHA), job safety analysis (JSA), training and qualification—all for the job that we imagine the workers will do. We *learn* about safety in the post-job review and operational learning. Both are essential, but one cannot be successful without the other. For too long we have placed our safety emphasis on better planning. We now know that real operational learning comes from understanding the difference between work as planned and work as done.

The problem with all this operational "play" in the work our workers do is the lack of knowledge your organization has about the work environment. As workers adapt and improvise solutions to "discovered work," new dangers are also being discovered. In the environment of discovery, our workers are now dealing with hazards that we have not planned for, mitigated for, protected for.

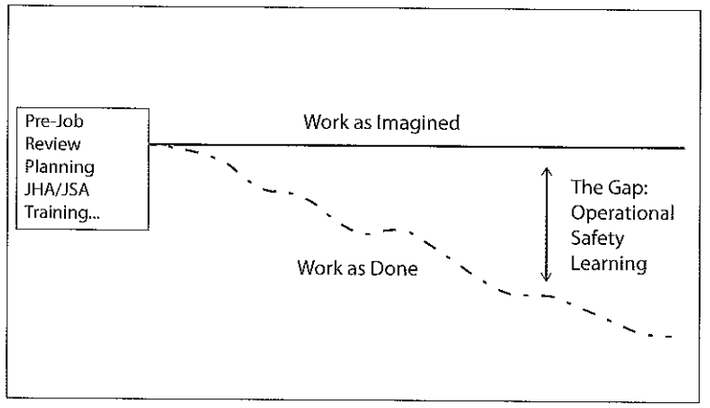

Figure 4.2 Operational Learning Happens Every Time a Task is Performed

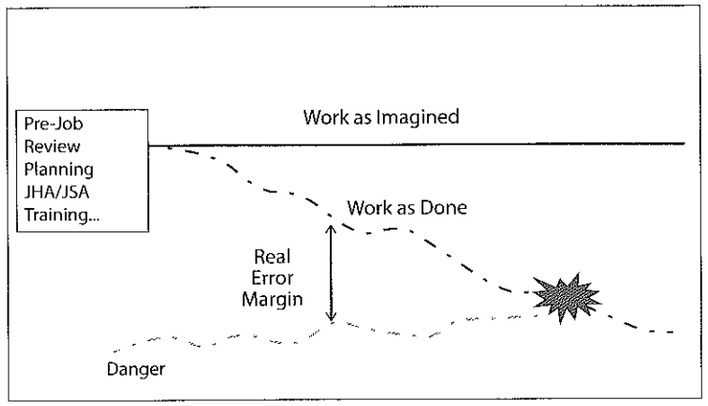

Figure 4.3 Actual Safety Margin

The best way to discover this information is to ask for it. You can do post-job reviews of successful work. Indeed, you *should* do post-job reviews of successful work. The only problem is that there is not much incentive to review successful work. We almost always are starting the next job before we complete the last job. That is simply our business reality. However, asking workers four quick questions which take little time or effort, and little to no money, will start to give you operational information, and reinforce a culture of "knowing" within your organization.

Tool: Successful Wrap Up

Four Post-Job Questions:

1. What happened the way you thought it would happen?
2. What surprised you?
3. What hazards did we identify and what hazards did we miss?
4. Where did you have to "make do," improvise or adapt?

Preprint these four questions on cards, the bottom of your job tickets, or as a part of a producton or manufacturing traveler ticket, and ask workers to spend a few seconds answering these questions. Pick these cards up. You are in charge of collecting and analyzing this information. Remember that once you have asked workers for the information, you will then have to do something with that information. They will expect to see an outcome or result.

The good news about the four post-job questions tool is that you will learn much about how work is done in your facility. You will be surprised by the

places in your processes and in your procedures where workers in the field, on their own, adaptively perform their work. The more you know about how work happens, the better prepared you will be to help workers create safety in practice.

Case Study
Seven Stitches, a Broken Finger, Cardinal Rules, and Termination

Sometimes all you need to do is ask the question, "how?"

NOTE

This case study is not only about the event, but also about the tremendous need to understand that sometimes safety culture change can take years, and sometimes safety culture change can happen quickly. In this case, the organizational shift in mindset happened swiftly and at just the right moment. This organization had just sent their senior management team to an 8 hour Human Performance training session. The managers loved, and understood, the class and the concepts, but after the excitement was over, the management was left wondering what they should do with this information. How would they apply the lessons?

It is not uncommon for people to think that these ideas make sense, but not see a way to apply these ideas to their organizations. It seems silly that the managers can't see how these ideas—the ideas of the new view of system performance—are mostly about how they choose to react to failure, but in the application of the ideas, blindness is often the case.

However, in this scenario, the safety people actually pushed the management team to ask, "how did this worker arrive in this position?" Nothing about this case made sense to the safety people. Pushing management to pause and take a deep breath before they pulled the termination trigger not only saved a worker his job, but also powerfully helped shift the culture of the organization. This outcome was a product of a courageous safety person explaining an event differently, more completely to his operational overlords

After this initial discussion, there was a substantial change in the organization. Suddenly, the managers stopped asking how to apply the new ideas they had heard in Human Performance class, and started applying them. Action in this case made all the difference. Your job is to determine whether the action in question was the event, or the shift in thinking—or both.

THE REPORT

Problem Statement: A senior maintenance technician received 7 stitches, and broke the tip of his right index finger.

The event transpired when a maintenance technician reached in and "jerked aggressively on the belt" of a moving conveyor lift. The technician had removed the guarding, and was able to grab the belt directly in front of the roller apparatus. Because of the technician's choice of hand placement, coupled with the fact that the belt was moving, the index finger of his right hand was yanked into the pulley apparatus, causing the break and the injury.

This event was investigated, and the root cause of the event was determined to be the worker showing poor judgment in reaching in to a moving machine: worker error.

Pursuant to the company's policy on cardinal rule violations (Rule 1: *Worker exhibiting undue or excessive risk*; Rule 2: *Working on an energized system without LOTO* [lock out, tag out]), it was recommended that termination proceedings should begin for this technician immediately.

This senior maintenance technician, who had years of practical experience and organizational history, was about to lose his job for replacing a conveyor lift belt that, as it turns out, had never been replaced before.

THE STORY

A conveyor lift is a stable performance tool in most production facilities. If this piece of equipment is actively scheduled and actively receives preventive maintenance, conveyors simply work—almost in spite of themselves. Conveyor lifts rarely break, and if maintained, can run for many years.

In this case, the conveyor lift moved glass bottles into a loader system hopper. This lift was old and finally starting to show some wear after many years of steady operation. Because this lift worked on elevating bottles and dumping them into a bottle hopper, this lift moved remarkably slowly—10′ a minute—much slower than the average speed a human walks. Conveyor belts that go this slowly last a long time and have little wear.

After operating on a production line 24 hours a day for over 20 years, this specific lift was finally starting to show some wear. The wear showing on the lift had indicated that it was time to replace the actual conveyor belt on this piece of equipment. The decision was made by the maintenance and production managers to replace the conveyor belt.

This specific belt had never been replaced, ever, by this manufacturing facility. The operating manual for this lift was secured. A new belt was ordered from the vendor, in accordance with both the belt numbers in the operating manual and the belt numbers available from the vendor. It was not difficult to order a replacement belt. The exact replacement, according to the numbers in the operating manual, was ordered and delivered to the manufacturing facility.

The maintenance on this machine was scheduled to begin on a Friday by the maintenance planners. It was also determined that because this specific belt moved so slowly, the belt replacement task would be done without stopping the production on the line. Instead, the decision was made to position a worker in this area of the production line, who would hand-carry empty bottles over the heads of the maintenance crews replacing and aligning the conveyor lift.

The Friday crew got the belt replaced, but this task took longer than was scheduled. This task had never before been done for this machine. Since this was the first time this belt had been replaced, the maintenance crew was learning to do the work while actually doing the work. In short, the job took substantially longer than anyone had imagined or anticipated it would take to complete. To make this situation even more complicated, the Friday crew was rotating, and would not be back at the plant on Saturday to complete the task. There would have to be a crew "handoff" on this job.

When the Saturday maintenance crew came in for their shift, they inherited a work task that was nearly complete. The new conveyor belt had been installed, but not tuned or aligned. At this point, although the belt was installed, it was not able to be operated until it had been seated and adjusted to run. To all intents and purposes, this should have been an easy job for this new crew. The production facility had other belts throughout the facility, though none ran this slowly. However, the process of adjusting a conveyor belt was not normally seen as difficult.

Belts are normally replaced with the machine off, the machine guarding removed, and with the power locked and tagged out. However, belt adjustments can only be performed with the belt operating. Most conveyor lifts have belt adjustment bolts that can be turned slightly to fine-tune a belt so that it operates in a straight line and doesn't work its way off the belt rollers. This was the case with this particular conveyor lift. Machine guarding was replaced on the belt rollers, the power was unlocked, and the machine was turned on in order to adjust the belts.

During this time, as was the case during this entire maintenance activity, the glass bottles were continually being hand-carried and dropped into the bottle hopper. Production had not been slowed or hindered for this maintenance activity. It is also necessary to note at this time the fact that this production facility was under heavy security because of what the company produced. Virtually every square inch of this facility was constantly monitored by video recording surveillance. Every second of this event was recorded and could (and would) be analyzed for a clear cause.

The Saturday crew was made up of a senior maintenance technician and two other lower level maintenance technicians—a normal distribution of skills

and talent. The senior maintenance technician in this case was well respected, with 20 plus years of experience in the facility, and was known as a safety leader. This crew was knowledgeable and capable, and normally needed very little supervision or direction.

The crew began the process of completing this conveyor belt replacement task. At around 9:00 am, the work began, and if everything had gone as planned it would have been done in a half hour or so. Two and a half hours later, the belt was not cooperating and was not yet installed.

When observing the video of this event, it is difficult not to notice the level of frustration of the workers increasing almost minute by minute. Nothing the senior technician did made any difference to the alignment of this belt. The video shows the technicians using a wrench to turn the set screws many times. Normally, these adjustments are small, but in this case, the crew could not get this belt to align correctly with both the end rollers and the tracking wheels. Both of these systems keep the belt reliably operating. During this time, the crew removed the machine guarding in order to assure the correct alignment of the belt.

At about the three-hour point in this process, the senior maintenance technician, moving up to the top of the belt—the section that was giving the crew the most trouble—grabbed the moving belt and jerked it with all of his strength. In the process of trying to force the belt into place, his hand was pulled between the belt and the belt rollers, breaking the tip of his finger, and lacerating his finger as his hand pinched out of the belt. No other worker, neither the other maintenance workers nor the worker tasked to manually carry bottles to the bottle hopper, attempted to stop this worker or even noticed his hand grab the belt.

THIS ACCIDENT WAS PREVENTABLE...OR SO THE ORGANIZATION WISHED...

The worker was injured. During the organization's normal, traditional event investigation, the decision was made that the cause of the event was the worker's poor choice in grabbing a moving belt running on an energized machine. If the worker had not grabbed this belt while the machine was running, this accident would not have occurred. If the guarding had been in place, it would have been impossible for the worker to stick his finger between the belt and the rollers, therefore this event would have never transpired.

In fact, both of those notions are true. Just as the following ideas are also true. If the belt had been the right size, the worker would not have become frustrated enough to try to force the belt into place, and this accident would not have happened. If production had been stopped, perhaps the crew would not have felt time and production pressure during this maintenance activity,

and with the pressure relieved, the accident might not have happened. If the belt had moved faster than 10' per minute, the technician would never have grabbed the belt, because a fast moving belt is a dangerous belt, and therefore the accident would not have happened.

The idea that by identifying the things that the worker should have done in order to avoid this accident the organization would be able to either change history or prevent the next event is foolish, but extremely attractive when answering the question, "why did this happen?" It is also a form of malpractice—by fixing only the worker, the company is allowing work to take place in ways that actually remove the workers' ability to control their work environment. Remember: operational surprise removes worker decision, and creates worker reaction.

Let's look at the conditions over which the worker had little or no control. I have made a partial list of the conditions that existed in the making of this failure. Notice how little control of the actual work environment the worker had in the context of this failure.

Conditions beyond Worker's Control:

- Maintained production status at facility
- Belt moved exceedingly slowly, giving the impression of manageable risk
- Belt had never been replaced since installation of equipment
- Belt could not be adjusted using adjustment screws
- Belt design and tolerances had changed
- Belt part number had not changed
- New belt would never fit in current machine configuration
- Belts must be adjusted while running
- No communication between maintenance crews due to schedule turnover

WHY THIS STORY IS DIFFERENT

With the new found knowledge of "another way to look at this event," the safety crew at this production facility started to try to understand the event from a much different point of view. The facility safety people knew that as soon as word of this event got to the corporate level, their choices would be significantly reduced. If corporate found out that this worker had stuck his hand in a belt that was running there would be three questions: Why was this device not locked out? Where was the machine guarding? Why haven't we used our cardinal rules to fire this guy?

So, the safety people made certain they had the story of this event *before* they called corporate to make their notification. Their belief was that if they could tell the story differently, more completely, perhaps they could cause a change in the course of the future.

It worked. By reporting the whole story to the powers that make decisions, the safety people ensured that the company's actions changed from the traditional "blame and punish" to a more "diagnose and treat" model. No one got fired. No one had to justify breaking a cardinal rule. The event didn't change—the way the safety people talked about the event changed, and that created a different conversation with the corporation.

How did the response to this event follow a different path? In short, the production facility made a conscious decision to respond differently to this failure. Even though this failure contained all the corporate triggers to immediate discipline of a worker, the trigger did not get pulled. The safety staff at this production facility presented a more complete contextual understanding of the failure. In the long term, they made the facility and its employees safer.

DISCUSSION

Some questions you should consider about this Case Study:

- Is there a need for "cardinal rules" in any organization's safety program?
- Who gets to draw the line that says a rule has been broken? Every rule is built with contextual and political flexibility—rules are always open to some interpretation.
- Does past adherence to a rule absolutely dictate future use of the same rule?
- Are there cases when a cardinal rule has forced the organization to make a less than optimal decision?

5
Change is Better When You Manage Change—And Change Needs to be Managed

QUESTIONS ABOUT YOUR ORGANIZATION:

- What changes have been made in your organization recently that have been especially successful?
- What made those changes successful?
- What did your organization do to make the changes stick?
- What did your managers do?
- What lessons can you take from other changes within your organization that you can replicate in order to help you change to your safety program?

MANAGING THE SAFETY CHANGE

You are now asking your management team to go on a journey. This is not an easy jump; the jump to a fundamentally different view of performance reliability does not occur naturally. You will serve as guide, teacher, enforcer, and diplomat for the new view and new operating parameters—and this job will never go away. You will be the conscience of the organization. You will be the leader and mentor. You will also be the follower and student. Just remember: this journey is well worth the effort. You will never know how much good you are doing for your organization. Just know that you are doing the best work you have ever done.

The biggest enemy of safety change is dogma – the belief that we already know the answers to the questions our organizations are forced to ask. If we know what the problems are, or similarly if we believe we already know the solutions, there is no need to go in to the field to learn more. There is no need to change. If we already are convinced our organization is doing the right things, just not doing those right things good enough, there really is no need to change. That is why our organizations keep doing the same corrections over and over without a different outcome. In order to change our organizations