Critical Facility: Power Supply Management – Case Studies

Contents

Critical Facility: Power Supply Management – Case Studies

Case Study 1: Critical Operation Failures

Module: Overview of Critical Operations

Topic: Facilities Operation Management and Critical Operations

Unfortunately, there are a lot of examples of Critical Operation failures throughout history that illustrate just how important it is to maintain the critical operations of Critical Operations facilities and the disastrous consequences that occur when systems fail. One such event was the partial nuclear meltdown at Three Mile Island in late March of 1979.

Three Mile Island, just outside of Harrisburg, Pennsylvania, was home to two working nuclear reactors that provided commercial nuclear energy to hundreds of thousands of American homes. On March 28, 1979, with the first unit powered down for refuelling purposes, the second unit experienced a fairly common blockage; but that blockage, in combination with numerous human-related and system-design errors, would eventually result in a full loss of the second unit and a near nuclear meltdown.

The accident started nearly a full 11 hours before the actual nuclear meltdown event in the secondary, non-nuclear system that supports the nuclear reactor itself. There was a blockage in one of the filters that cleans the water used to cool the system, which operators used compressed air in an attempt to unblock. However, they were unaware that a check valve within the system had become stuck open because a key light on the system's control panel falsely indicated that the valve was closed; when the compressed air was used to clear the blockage, the cooling water made its way past the stuck-open valve and into an instrument air line. This would eventually cause the pumps that fed water into the steam generator that fed the cooling system to turn off hours later, which increased heat and pressure in the reactor coolant system and caused the unit to perform an emergency shutdown.

But the problems didn't end there: when the emergency shutdown occurred, control rods were automatically inserted into the nuclear core to halt the chain reaction, but decay heat was still being created and wasn't being removed from the system because of the unknown issues in the water loop. Auxiliary pumps were automatically activated by the process to pump water into the system, but the valves used within the pumps had been closed for routine maintenance—an action that was a flagrant violation of key regulatory rules. The increasing heat and pressure, and the inability for the secondary system to intervene, caused the automatic opening of a relief valve, which should have closed once the heat and pressure had been released. But, of course, the relief valve was experiencing its own mechanical failure and was also stuck open, allowing coolant water to enter the nuclear reactor itself, and resulting in the partial nuclear meltdown of Unit 2.

Once the partial meltdown occurred, it was possible that radioactive materials were released into the environment. While some reactive gases were definitely released, in the initial investigations after the accident, the plant ownership deemed it unlikely that any dangerous levels of radiation had been allowed to reach the community. However, residents were urged to stay indoors and, as a precaution, the Governor called for a voluntary evacuation of the local population—though all of these emergency efforts took place *days* after the actual event. Unit 2 was subsequently decommissioned and remains permanently shut down to this day, while Unit 1 continues to generate nuclear power.

Critical Facility: Power Supply Management – Case Studies

What was most troubling about the Three Mile Island accident were the system design errors and human errors that allowed such an event to occur in the first place. The light that signalled the stuck-open valve should have indicated the problem. But, beyond that, the operators should have realized that the continuing problems were indicative of a larger problem, rather than going on the assumption that the faulty light was correct and allowing hours to go by before a different shift of operators came on duty and were able to correctly diagnose the problem.

Fortunately, a complete Critical Operation failure of this nature could have resulted in far worse consequences, including major casualties to both the general population and the environment. But what was learned from the Three Mile Island accident was just how important it is to have a complete understanding of the various components of a mission critical system and the need for operators to be aware of the interdependence of these components—and how a failure of one component can have devastating effects on the rest of the system

.

Critical Facility: Power Supply Management – Case Studies

Case Study 2: IT Failures

Module: Overview of Critical Operations

Topic: Facilities Operation Management and Critical Operations

IT failures can have disastrous effects on the day-to-day operations of critical organizations. Take, for instance, the three-hour outage in 2013 that brought the NASDAQ stock exchange to a screeching halt at the height of trading hours.

Trades on the various stock exchanges are conducted primarily via data feeds between the respective exchanges' platforms. A reliable, consistent connection between them is essential for trading to occur; essentially, NASDAQ's day-to-day operations is entirely reliant on the stability of its IT infrastructure. Unfortunately, a number of IT-related issues would be to blame for the hours-long

outage that prevented trading communications between the various platforms that process quotes and trades for the stock exchanges.

It started with a flood of data from the New York Stock Exchange's Arca platform to NASDAQ's Securities Information Processor (SIP), which swamped SIP's capabilities. Unfortunately, this massive data dump uncovered a fatal flaw in SIP's software code: the stream of data overpowered the system, rendering it incapable of triggering its backup system and, as it was unable to failover, there was an outage. Because these two platforms communicate quotes and trade data between one another, trading essentially ground to a halt for the duration of the outage. Without a backup system in place to handle transactions and processing during the outage, trades had to come to a complete stop. While it only took about 30 minutes for the data feeds to be repaired and fully operational, it took more than three hours of testing and evaluation before actual trading could resume.

While it is unknown how many trades the outage affected or prevented, you can do some simple math to estimate its effect. Data taken from a few years before the outage shows that the NASDAQ SIP processed an average of about 6.25 million trades per day; that equates to about 1 million trades per hour during open market hours. So, for the three-plus hours that the system was experiencing its outage, more than three million trades were unable to be processed. In the grand scheme of things, the NASDAQ outage might not have been a catastrophic event in the sense that there were no casualties or major losses, but the disruption of more than three million trades over the span of just three hours is a significant blow to a nationwide financial system that relies heavily on the stability of its IT infrastructure.

Critical Facility: Power Supply Management – Case Studies

Case Study 3: Security Breaches

Module: Overview of Critical Operations

Topic: Facilities Operation Management and Critical Operations

Security breaches today are most often discussed in relation to data security breaches or the unwanted access of secure information. One of the more memorable data breaches of recent memory was the hack of retail store Target's databases in late 2013, during which the debit or credit card information of more than 40 million customers and the personally identifiable information (PII) of more than 70 million customers was accessed.

The breach wasn't particularly innovative or aggressive: in late November of 2013, malware was installed on the company's security and payment systems, which was designed to capture credit card numbers during checkout and then store them on a specific server within the same network that was being controlled and accessed freely by the breaching parties. From there, the data was shuffled to

multiple servers around the country as staging points, and then moved to computers in Russia where the hackers could use the credit card numbers to make unauthorized purchases.

However, Target had already installed malware detection tools on the servers used for the security and payment systems, and the security specialists who monitored them were aware of the suspicious activity almost as soon as it started. They notified the company's security operations center of the potential breach but, unfortunately, these alerts were either not taken seriously or weren't acted upon with immediacy. By mid-December, when the company disclosed the breach to the general public, more than 40 million customer debit and credit card numbers and 70 million customer names, addresses, phone numbers, and email addresses had been stolen.

Since the breach, close to 100 lawsuits have been filed against the company on behalf of both customers and financial institutions, claiming negligence on Target's part and seeking compensation for losses sustained as a result of the breach. Between these lawsuits and its customer response effort, Target has spent millions of dollars in response to the data breach.

Critical Facility: Power Supply Management – Case Studies

Case Study 4: Failing to Adhere to Procedures or Policies

Module: Overview of Critical Operations

Topic: Facilities Operation Management and Critical Operations

Policies and procedures exist in any organization for a purpose: to ensure the safe, consistent operations of key systems or devices. So, when they aren't followed or adhered to strictly, things often go wrong—sometimes with devastating consequences. One such case of the disastrous effects of failing to adhere to policies and procedures is the nuclear accident at the Chernobyl power plant in 1986.

The Chernobyl power plant was home to four operating power reactors and a fifth reactor under construction, on a site just north of what is now Kiev, Ukraine (at the time, it was still part of the Soviet Union). The reactors were a Russian design called RBMK, which was fueled by natural uranium, cooled by water, and moderated using graphite construction. However, this design was considered "fundamentally faulty" and "having a built-in instability" by members of the expert community, especially given the fact that, rather than shut itself down in the event that is starts to lose coolant, an RBMK actually would increase in reactivity and begin to run hotter and faster. Additionally, the reactors were not protected by containment structures, which were required of reactors in other countries. From the start, the nuclear reactors on the Chernobyl site were already not in compliance with the standards expected of a safely operating nuclear facility.

Unfortunately, on top of a poorly designed system, the operators at Chernobyl were fundamentally lacking in the knowledge, skills, training, and awareness of the magnitude of their actions needed to safely operate the facility. In fact, it is without a doubt that the accident at Chernobyl was a direct result of actions taken on the part of the operators that failed to adhere to any proper protocol.

The immediate cause of the accident was "mismanaged electrical engineering equipment," but the larger story was that a number of engineers—with no working knowledge of the physics of nuclear reactors—had decided to run an experiment to see if they could use the rotational energy from the steam turbine in the reactor to generate additional backup power for the plant. Before performing the test, they disconnected every important safety system, including the emergency core-cooling system, and disconnected every backup electrical system—all the systems that would be needed in the event of an emergency to help them safely operate the reactor controls.

The experiment went poorly from the start: the operators reduced one reactor's power level to wind up the turbine, but did so far too quickly, which caused a build-up of fission by-products in the reactor core and poisoned the reaction. In an effort to stabilize it, the operators began to remove the control rods from the unit. Despite their efforts, they couldn't increase the power level to more than 30 megawatts, an operating level in which the instability of the reactor is at its potential worst and a level that was specifically forbidden by the plant's own safety rules. They continued to remove more control rods until only 6 of more than 210 rods remained, even though the minimum number of control rods for the proper documented operations of the reactor was 30.

But the engineers pushed on with their experiment, next shutting down the turbine generator, which reduced power to the water pumps and thus reduced the flow of coolant water to the system. This unfortunately increased the chain reaction in the reactive materials, and the power level in the reactor surged. While there is some debate over whether it was an operator decision or an automatic system action, the control rods that had been removed were hurriedly re-inserted into

the reactor. Unfortunately, the control rods had a fatal design flaw: their graphite tips increased the reaction, and displaced water from the rod channels increased it further. There was simply too much reactivity, and the reactor exploded.

Between the steam in the initial explosion and the resulting fires, at least five percent of the radioactive materials from the reactor were released into the air and travelled downwind. Two plant workers died in the explosion and another 28 emergency responders died in the following weeks due to acute radiation poisoning. Containment and clean-up of the site took months, involving more than 500,000 workers (many of whom were exposed to dangerous levels of radiation) and costing more than 18 billion roubles (or more than 3 million U.S. dollars, using the current exchange rate). But the impact of the accident has been felt for decades. Since the accident, more than 230 people have suffered acute radiation sickness, and increased incidences of cancers and other health disorders likely to be related to the accident have been reported.

Critical Facility: Power Supply Management – Case Studies

Case Study 5: Improper Documentation

Module: Overview of Critical Operations

Topic: Facilities Operation Management and Critical Operations

As you now know, having access to documentation for the design, proper operation, and even the individual components of Critical Operations is extremely important for the consistent, safe, and reliable functioning of a facility. And when that documentation is lacking or non-existent, it can be difficult to maintain the systems' optimal functionality. Case in point: the improper documentation for a new data center supporting a large research university.

Typically, the university would build on their own property and construct all facilities to their standards and then maintain and operate the facility after its construction. But this time around, the design and construction of the facility was provided in what is known in the industry as a turn-key approach—it was designed and constructed by a company that specializes in this type of facility, to their own standards and specifications, with input from the university's facilities and IT staff only regarding its specific space and technical needs. The data center was built on a leased site five miles from the center of campus and was operated by staff contracted by the specialty company, not by university IT staff.

On a particularly hot summer day, the temperature in the primary server room in the data center began to rise. Fortunately, the protocol to address this issue was documented in both hard copy and electronic form, so the operators knew where to look. Despite their best efforts to adjust the system according to the documentation, it became clear that the system wasn’t responding appropriately. Clearly, they were not addressing the correct issue, but their documentation didn't have any other potential causes or possible solutions. The operators called upon the university’s facilities staff to help them, but since they had not been directly involved in the design or construction of the facility, they weren't able to provide the operators with any helpful information.

After some time and exploratory research, it became apparent that the temperature rise was the result of an equipment failure, specifically a malfunctioning relay panel. Unfortunately, the data center operators didn't have any spare components on hand to fix the issue immediately. While the university facilities and IT staff typically had spare parts, the data center used very specific equipment used by the specialty company, and they didn't have a compatible component. Instead, the data center operators had to contact the firm that designed and constructed the building for the necessary parts, which would be accompanied by specific installation instructions. However, the component wouldn't be delivered for another 24 hours, and in the meantime, the temperature in the server room had continued to rise and was getting dangerously close to hazardous levels. Fortunately, the design and build firm was also able to talk the operators through some specific adjustments to the system that would temporarily provide a solution until the component was replaced, and after a few hours of trial and error, the server room was back to safe operating conditions.

Critical Facility: Power Supply Management – Case Studies

Case Study 6: Power Loss and Preparation

Module: Power and Power Sources

Topic: Power Fundamentals

Suffering a power loss is one of the worst things that can happen to MCOs, as it will bring down the entire operations of a facility that is critical to keep up and running. One of the most unpredictable events that can lead to widespread power outages are weather-related natural disasters. Case in point: the devastating effects from the havoc that Hurricane Irene wreaked on some of the most critical facilities along the eastern seaboard, including power loss at one major national telecom company's MCOs.

On a Friday in late August 2011, facility managers and CRITICAL OPERATION personnel across the country shifted into high gear as Hurricane Irene—a massive category 1 storm—was about to make landfall in the Carolinas. While proper preparations were in place to the extent possible, not everyone was ready to handle the $14 billion in damage Irene would leave behind after driving straight up the eastern seaboard through New England.

One team of Critical Facility Managers running the national telecom company's Critical Operations started response activities even before the first power lines went down. In the days or hours before a storm like this, they would begin inventorying fuel levels at all the CRITICAL OPERATION locations and calculate available run times should outages occur. Knowing they were in for a long weekend, the team also began aggregating vendor contact information so other FMs could take over response efforts while eastern regional FMs could catch rest when possible. In this case, as the forecasts sharpened up and it was evident there would be widespread power outages across many states, additional fuel supply vendors were brought on board.

Once power went down at the MCOs, the FMs ran through checklists: contacting client managers to check the status of supported technology, verifying with any onsite personnel that backup equipment had picked up the critical power load and generators were starting, and having site walks done to check the infrastructure. All told, 17 sites lost power, 2 for as long as 4 days. All but one site sailed through fine with no interruptions to Critical Operations except for some minor equipment failures (backup pumps developing minor leaks, for instance) that could be dealt with after the storm.

The one site that went down did so because the output breaker of the generator failed. It was an older site (the breaker itself was probably at least 15 years old), so the equipment failure was not that shocking. The problem, though, was that the customer had previously refused to follow the team's recommendation to update the building's load transfers as part of the electrical maintenance program. It had been years since that breaker had been cycled, so who knows: perhaps that failure could have been found during a controlled, scheduled maintenance evolution and repaired, instead of coming to light during the height of the storm and causing a critical outage.