PATHFINDER REPORT

Safeguarding your Critical IT Workloads

Operational Excellence in Multi-Tenant Datacenters

DECEMBER 2016

COMMISSIONED BY

About this paper

A Pathfinder paper navigates decision-makers through the issues surrounding a specific technology or business case, explores the business value of adoption, and recommends the range of considerations and concrete next steps in the decision-making process.

About 451 Research

451 Research is a preeminent information technology research and advisory company. With a core focus on technology innovation and market disruption, we provide essential insight for leaders of the digital economy. More than 100 analysts and consultants deliver that insight via syndicated research, advisory services and live events to over 1,000 client organizations in North America, Europe and around the world. Founded in 2000 and headquartered in New York, 451 Research is a division of The 451 Group.

NEW YORK
20 West 37th Street
New York, NY 10018
+1 212 505 3030

SAN FRANCISCO
140 Geary Street
San Francisco, CA 94108
+1 415 989 1555

LONDON
Paxton House
(Ground floor)
30, Artillery Lane
London, E1 7LS, UK
P +44 (0)207.426.1050

BOSTON
One Liberty Square
Boston, MA 02109
+1 617 598 7200

I. Executive Summary

Enterprises, government agencies and SMBs globally are increasingly turning to an outsourced datacenter business model that leverages colocation providers, also known as multi-tenant datacenter (MTDC) providers. When customers consider MTDC options, they typically look at a number of attributes and requirements – including facility location; cabinet and power costs; facility quality and reliability; connectivity options and costs; power and cooling capacity and densities supported; energy efficiency; security; fire detection/suppression; expansion capacity; and services offered, in addition to SLAs. However, there is one attribute that we believe more customers should also look for – what we call operational excellence.

Operating a datacenter is complicated business, and best left to experts. Datacenters are large and complex systems that are becoming even more complex to enable support for modern IT equipment with higher power densities. Many organizations where datacenters are not a core competency will not have the skillset necessary for ongoing and effective datacenter management and operations. However, as business reliance on IT systems located within datacenters continues to expand, the importance of datacenter operations has increased proportionally, in terms of ensuring datacenter resources are performing optimally and are continuously available.

Datacenter outages are incredibly costly, and typically happen for mundane, preventable reasons. Datacenter outages are also becoming more costly each year, particularly when the value of lost revenue and expense of business disruption is factored in. Yet many of the most costly and visible datacenter outages in the past year could have been prevented if proven sets of standards and processes had been established and followed. Standardizing on a datacenter operations paradigm of continuous improvement and excellence is critical in order to identify weak areas and address problems prior to them becoming crises that result in datacenter downtime. There is always opportunity for improved operational capabilities in complex datacenter environments.

The core business of MTDC providers is offering datacenter services to customers, and a critical part of this business is operating datacenters for customers with the highest possible availability and uptime. So operating datacenters is indeed a core competency for MTDC providers. By driving standardization, accountability and transparency in the reporting chain for datacenter operations, service providers can also focus on innovation and customer satisfaction.

II. Datacenter Operations Complexities

Enterprises have needed space for their IT equipment since they've had IT equipment. Until about 15 years ago, servers and other IT hardware were almost always kept in closets or larger datacenter spaces that enterprises had constructed, typically in or near their offices. These facilities are aging and, in many cases, are not aging well. Datacenters and server rooms built during the 1980s and 1990s, even with state-of-the-art systems at the time, are in many cases failing to hold up to the latest IT requirements. The complexity of datacenters has also grown considerably in recent years due to increases in equipment power density and the resulting power and cooling requirements, and a decreased tolerance for outages. At the same time, enterprises are also being pressured to increase IT redundancy and reliability due to increased regulation, greater internal reliance on IT services and the vastly greater use of the Internet to communicate with and sell to customers and partners. In addition, for a good many businesses, building and operating datacenter space is not a core competency or a competitive differentiator. So while it may be easy to think that operating a modern datacenter is in some way analogous to running a server closet, and that no extra skillsets would be required, this is increasingly not the case.

Datacenters are large and complex systems that can be subject to failures — sometimes catastrophic or even potentially fatal failures. High-voltage electrical systems, large-scale mechanical and infrastructure components, high-pressure water piping, power generators, and other sophisticated elements all combine to create the complex systems necessary in modern datacenters to sustain today's IT loads. The greater the number of components and the higher the energy and heat levels, velocity, size and weight of these components, the greater the skill and teamwork required to plan, manage and safely operate the systems within datacenters. Between mechanical components and human actions, there are *thousands* of possible points where an error can occur and potentially trigger a chain of failures. As the datacenter environment becomes ever more complex, more effort is required to maintain it: More systems are required to monitor it, and more processes and procedures must be defined and followed to reduce the chance of error or failure. Complexity of datacenter systems breeds complexity of datacenter operations.

Complex systems science suggests that most large-scale complex systems, even well-run ones, by their very nature are operating in 'degraded mode' – i.e., close to the critical failure point. This is due to the progression over time of various factors, including steadily increasing load demand, engineering forces and economic factors. It is a truism that complex systems

tend to fail in complex ways. Again and again, we see that it is not a single factor but the compound effect of multiple factors that disrupts sophisticated systems found in datacenters. Often referred to as 'cascading failures,' complex system breakdowns usually begin when one component or element of the system fails, requiring nearby 'nodes' (or other components in the system network) to take up the workload of the failed component. If this increased load is too great, it can cause other nodes to overload and fail as well, creating a waterfall effect as every component failure increases the load on the other, already stressed components.

A common trend being seen in datacenters is higher power densities per cabinet. Server density is driven by a mixture of engineering forces, such as advancements in server design and efficiency, and economic pressures, as well as demand for more processing capacity without increasing the facility footprint. Maximizing capacity, increasing density and hastening production from installed infrastructure improves the return on investment (ROI) on these major expenditures. Increased density also results in increased numbers and complexity of critical electrical and cooling elements. Now the system is running at higher risk, with more components, each of which is subject to individual fault/failure, as well as more power flowing through the facility, more heat generated, etc. The enormous investments in datacenters and other highly available infrastructure systems create conditions of elevated risk and higher likelihood of failure. Further exacerbating the problem, deferred maintenance, whether due to lack of budget or hands-off periods due to heightened production, further pushes equipment toward performance limits – the breaking point. Whether high density or not, the potential for a catastrophic outcome is inherent in the very nature of complex systems with this dynamic mix of forces. For large-scale mission-critical and business-critical systems, the implication is that designers, system planners and operators must acknowledge the potential for failure and build in safeguards.

Large systems such as power plants, airplanes and oil rigs and the industries that use them build in safeguards against failure, and as such they employ multiple layers of protection and backup. To safeguard against failures, standards and practices have evolved in many industries that encompass strict criteria and requirements for the design and operation of systems, often including inspection regimens and certifications. Compiled, codified and enforced by agencies and entities in each industry, these programs and requirements help protect the service user from the bodily injuries or financial effects of failures, and spur industries to maintain preparedness and best practices.

As large complex systems, datacenters require similar establishment of codified operational practices and standards. Based on analysis of over 20 years of datacenter incidents, the Uptime Institute, 451 Research's sister company, has concluded that human error in datacenters signifies management failure to drive change and improvement, and finds only single-digit percentages of spontaneous equipment failures in datacenters among overall incidents. Leadership decisions and priorities that result in a lack of adequate staffing and training, a lack of well-defined processes and procedures to manage and operate complex datacenter systems, an organizational culture that becomes dominated by a fire-drill mentality, or budget cutting that reduces preventive/proactive maintenance can result in cascading failures.

III. Impact of IT Outages and Root Causes

Datacenter failures are not new – in fact, per hour of aggregated uptime, they are far less common than they use to be, thanks in part to new technologies. The instrumented datacenter is saving itself more and more, and equipment redundancy is responsible for more saves every year. This year has been no different. Back in 2014, 45% of incident 'saves' were due to equipment redundancy, while this year that figure is 58%. Meanwhile, technician intervention went down from 32% to 25% this year. However, increased dependency on IT and increased codependency between IT systems means the impact of failures reverberates ever more widely. 2016 has seen some high-profile datacenter incidents that made national news. Of these, the downtime suffered at Delta and Southwest Airlines are the most notable. Unlike in 2014 and 2015 – when security breaches caused the biggest problems (at Sony and UK telecom provider TalkTalk, for example), there are some common themes in the most recent failures seen in Figure 1. Failures in power distribution equipment have been the root cause of several incidents, and problems with IT recovery have often amplified the severity of the issue.

Figure 1: High-Profile Datacenter Failures – 2H 2016

Source: 451 Research, 2016

Company/datacenter(s)	Date (s)	Affected areas/ extent	Cause	Cost?
Delta Airlines	8-Aug	All operational systems in NA.	Power surge, power/ transfer switching failure; IT systems corrupted. Some servers didn't have dual power chords?	1800 flights cancelled. Quarterly earnings expected down 10%.
Southwest Airlines	20-Jul	All operational systems in NA. 12 hour outage, cancellations for several days.	Malfunctioning router triggered multiple problems (IT level).	2,300 flights cancelled. Estimated loss of $177m in passenger revenue
TeleCity LD8 (Equinix)	19-Jul	Some Linx traffic. BT Broadand.	UPS failure	Not known/undisclosed
Telehouse	21-Jul	UK and beyond. BT Broadband/email services in UK. 7-10 hours.	"Tripped circuit breaker".	Not known/undisclosed
Telehouse	21-Jul	UK and beyond. BT Broadband/email services in UK. 7-10 hours.	"Tripped circuit breaker".	Not known/undisclosed
FCA @ Fujitsu Sunnyvale CA	24-27 Sep	System for managing 50,000 FCAs.	Transformer failure?	50K financial institutions unable to access. Strategically embarrassing.
ING Bucharest	10-Sep	Banking systems.	Noise from fire suppression systems damages dozens of disk drives.	Systems down for 10 hours. Many storage systems and servers replaced.
SSP at Solihull datacenter.	26-Aug – 24-Sep (?)	All core systems.	Power outage at Solihull triggered SAN problems. Second SAN failure followed. Attempting emergency migration to Tier 3.	40% of UK insurance brokers unable to access renewals data.
Global Switch 2, London	10-Sep	Many customers affected, notably Claranet.	222ms high voltage drop/ circuit breaker/DRUPS caused 222ms break, triggering shutdowns. Claimed Tier 3 standards...	Not known/undisclosed
Global Switch 2, London	6-Jun	Many customers affected.	Lightning strike led to several hours outage for some customers.	Not known/undisclosed

Regardless of size, a big portion of both the operating and capital budgets of datacenters is typically dedicated to maintaining availability in the form of redundant equipment and rigorous processes. This expenditure and effort are considered 'table stakes' – the risks and costs of failures are so high that most businesses opt for a very high level of resiliency with little cost-benefit analysis. However, a recent study published by the independent Ponemon Institute, sponsored by Emerson Network Power, shows that datacenter downtime is in fact very costly, and becoming even more so. The study's findings suggest that there is still a strong need for continued investment in site-level availability.

The '2016 Costs of Data Center Outages' study – the third in a series published by the independent Ponemon Institute, sponsored by Emerson Network Power – provides some useful financial information and insights. The institute's research into the causes and costs of datacenter outages is based on unplanned failures at 63 datacenters in the US, representing 15 industry segments, during a 12-month period. Costs were calculated using an activity-based costing model, with cost centers (activities) as follows: detection; containment; recovery; ex-post response (after-the-fact incidental costs); equipment; IT productivity loss; user productivity loss; and third party (use of contractors, advisors required to recover). The researchers also added two large consequential costs: lost revenue and business disruption (such as damage to reputation, customer churn and lost business opportunities).

Using data from previous reports, the study found that the average (mean) cost of datacenter downtime is increasing:

Figure 2: Datacenter Outage Costs and Length

Source: Ponemon Institute, 2016

YEAR	AVERAGE COST (USD)	AVERAGE DURATION
2010	$505,502	97 minutes
2013	$690,204	86 minutes
2016	$740,357	95 minutes

These averages include partial and total shutdowns. Partial shutdowns in the study tend to last less than half as long as total shutdowns – most likely an indication of severity. There is a linear relationship between the length of downtime and the cost – an argument in favor of the use of generators and recovery sites. What the report doesn't discuss in detail is why the costs of downtime are rising. However, it seems reasonable to assume that, inflation aside, the level of dependency on online systems is increasing every year, and, equally, that improvements in IT systems and utilization mean that more IT work is supported per kilowatt of datacenter capacity than in earlier years. Unsurprisingly, the biggest costs associated with failures were not in direct areas such as recovery or equipment replacement, but in consequential losses – notably business disruption and lost revenue. Overall, these indirect and opportunity costs accounted for 61% of losses. The study also found that IT failures and cybercrime failures cost the most. We speculate this is because these failures may take more time to analyze and repair, leading to greater consequential losses. However, there are very few IT-system-related failures – perhaps because if they are critical, they have their own system-level resiliency built in.

The cost of downtime to any organization dependent on IT is potentially huge, and this invariably justifies the enormous effort and expense that is invested in this area. For datacenter owners and operators, the question is becoming increasingly complicated: new architectures and software-based models may – in some situations – make it possible to reduce the level of physical redundancy and make more use of remote services or software/data failovers. Ponemon's study highlights two main facts: first, the costs of datacenter downtime are as high as ever and rising, and second, many of the failures can be prevented by investment in processes as well as equipment at a physical level.

ROOT CAUSES OF DATACENTER OUTAGES

The Uptime Institute has done extensive research on the causes of datacenter incidents through the Uptime Institute Network's Abnormal Incident Reports (AIRs). Each network member company has the ability to anonymously submit AIR data, which is then searchable by other members (with the submitting company's name omitted). Member companies use this as a tool to not only research incident types and causes, but also to identify issues with different brands and types of equipment. The data has been gathered from a broad set of industries, including financial and business services, telecom, technology, healthcare and others.

It is important to study datacenter incidents as they involve many types of equipment found in datacenters with thousands of possible points where an error can occur, which demonstrates the complexity of overall datacenter operations. Each of these many types of equipment must be constantly monitored and maintained for reliable datacenter operations without impacting the uptime of IT equipment in the datacenter. One important distinction to note is that a datacenter incident does not necessarily mean an outage; in fact it rarely does. Availability failures account for less than 10% of all incidents reported (though, frankly, it is at least possible that some companies only report incidents that were 'saved,' due to some perceived risk in the reporting process). As discussed, however, these incidents bring with them the higher probability of kicking off some sort of cascading event, where human error tends to play a much bigger role in an ultimate failure, whether it be an actual mistake in handling the event, or overall mismanagement.

Figure 3: Abnormal Incident Reports (AIRs), 1995-2015 (n=6,297)

Source: Uptime Institute, 2016

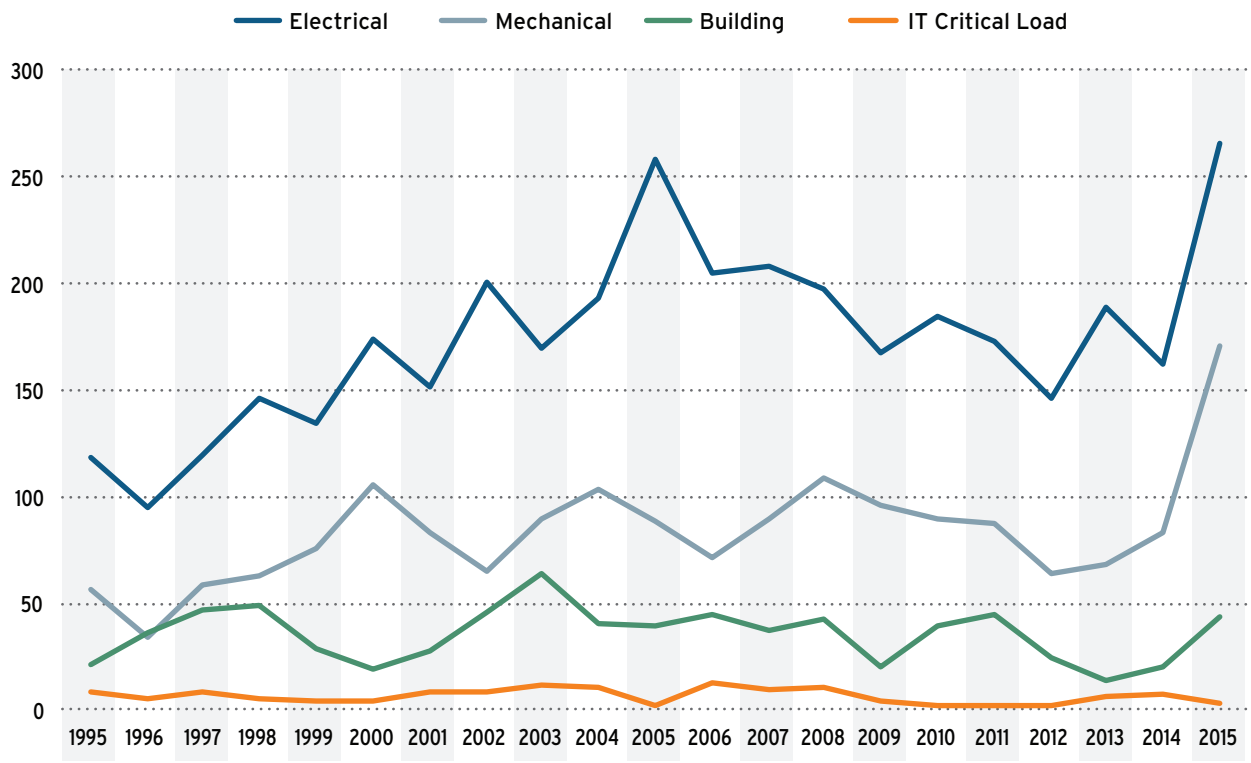

For the purposes of this paper, we break AIRs down by division (the broad categories of electrical, mechanical, building and IT critical load). Looking at 21 years' worth of AIR data in Figure 3, the electrical and mechanical divisions are associated with more incidents by far; electrical even more so than mechanical, typically due to its on/off nature. When looking at this data, though, it's an important distinction to note that this equipment was *associated with* the incident, because it may not actually have been the cause. It's also important to note with this particular chart that the jump in incidents in 2015 is due simply to a 90% increase in reporting over previous years, and is not necessarily representative of an industrywide increase in incidents.

Figure 4: Electrical Division AIRs, 1994-2015

Source: Uptime Institute, 2016

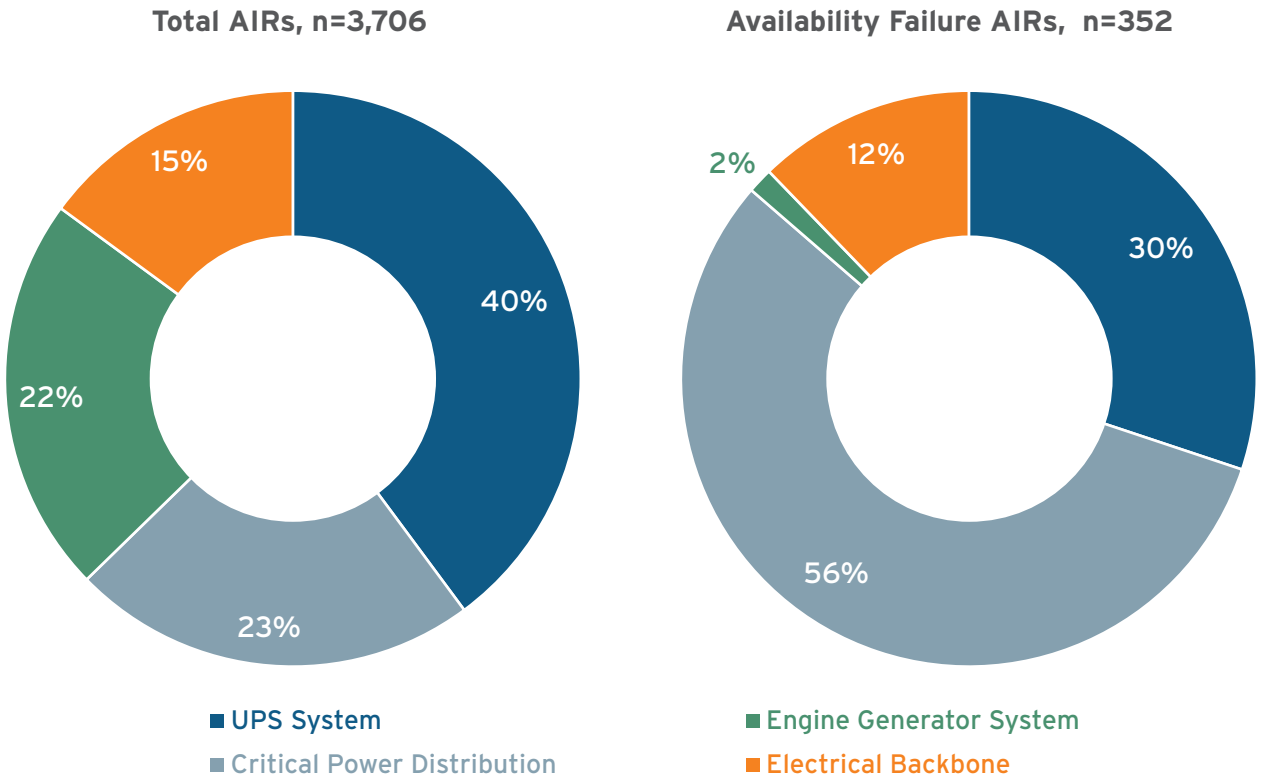

Starting first with electrical division AIRs in Figure 4, UPS is the leading system linked to electrical AIRs. However, when viewing this data by actual outages vs. non-outages, the story changes quite significantly. In terms of availability failures, the culprit is usually critical power distribution, rather than UPS systems.

Figure 5: Mechanical Division AIRs, 1994-2015

Source: Uptime Institute, 2016

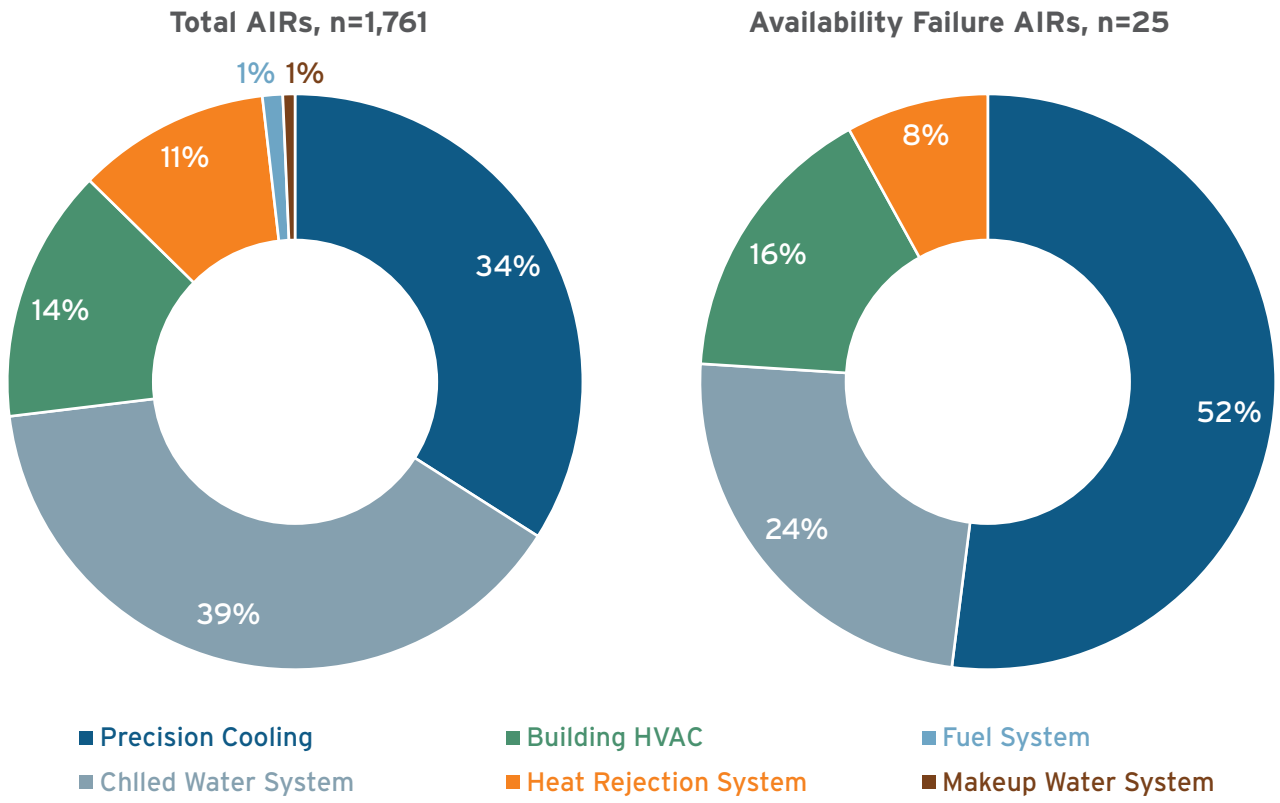

When looking at the mechanical division in Figure 5, we see a similar narrative. At a high level for all mechanical AIRs (outages and non-outages), we see they are primarily due to the chilled water system and precision cooling (CRAC, CRAH, enclosed systems, etc.) over the past 21 years. But, while not evident in Figure 5, in 2014-2015, building HVAC had increased AIRs – this includes air-handling units, makeup air or relief exhaust, and pump HVAC – while precision cooling has had fewer AIRs associated with it recently than pre-2013. So precision cooling is actually on its way down as a share of the total, and hopefully this trend continues. Chilled water systems, on the other hand, are consistently associated with a high number of AIRs. However, when we consider outages versus non-outages, precision cooling still sticks out like a sore thumb as a major factor.

If we dig further into the data and look at actual root causes, rather than just the equipment *associated with* the incidents, we see that generally it is a lack of maintenance that causes incidents. It is worth noting that while preventative maintenance accounts for roughly 28% of the overall 'human error' segment, this maintenance has never been responsible for an outage – at least as far as the AIRs data is concerned. This does, however, lead to many questions around what exactly constitutes human error. We've already outlined the philosophy that most outages (not just incidents) within these complex systems are due to human error, and two recent incidents demonstrate how the dynamics of complex systems failures can quickly play out in the datacenter environment.

Example A

A Tier III Concurrent Maintenance datacenter criteria requires multiple, diverse independent distribution paths serving all critical equipment to allow maintenance activity without impacting critical load. The datacenter in this example had been designed appropriately with fuel pumps and engine-generator controls powered from multiple circuit panels. When built, however, a single panel powered both, whether due to implementation oversight or cost-reduction measures. The issue is not the installer, but rather the quality of communications from the implementation team and the operations team.

In the course of operations, technicians had to shut off utility power to an electrical switchgear during the performance of routine maintenance. This meant the building was running on engine-generator sets. However, when the engine-generator sets began to surge due to a clogged fuel line, the UPS automatically switched the facility to battery power. The day tanks for the engine-generator sets were starting to run dry. If quick-thinking operators had not discovered the fuel-pump issue in time, there would have been an outage to the entire facility: a cascade of events leading down a rapid pathway from simple routine maintenance activity to complete system failure.

Example B

A Tier IV Fault Tolerant datacenter criteria requires the ability to detect and isolate a fault while maintaining capacity to handle critical load. In this example, a Tier IV enterprise datacenter shared space with corporate offices in the same building, with a single chilled water plant used to cool both sides of the building. The office air-handling units also brought in outside air to reduce cooling costs.

One night, the site experienced particularly cold temperatures, and the control system did not switch from outside air to chilled water for office building cooling, which affected datacenter cooling as well. The freeze stat (a temperature-sensing device that monitors a heat exchanger to prevent its coils from freezing) failed to trip; thus, the temperature continued to drop and the cooling coil froze and burst, leaking chilled water onto the floor of the datacenter. There was a limited leak-detection system in place and connected, but it had not been fully tested yet. Chilled water continued to leak until pressure dropped and then the chilled water machines started to spin offline in response. Once the chilled water machines went offline neither the office building nor datacenter had active cooling.

At this point, despite the extreme outside cold, temperatures in the data hall rose through the night. As a result of the elevated indoor temperature conditions, the facility experienced myriad device-level (e.g., servers, disc drives and fans) failures over the following several weeks. Although a critical shutdown was not the issue, damage to components and systems – and the cost of cleanup and replacement parts and labor— were significant. One single initiating factor – a cold night – combined with other elements in a cascade of failures.

In both of these cases, relying on front-line operators to save the situation is neither robust nor reliable, and the fingerprints of human error can be seen in both examples. In Example A, the electrical panel was not set up as originally designed, and in Example B the leak-detection system, which could have alerted operators to the problem, had not been fully activated. Facility infrastructure is only one component of failure prevention; how a facility is run and operated on a day-to-day basis is equally critical.

Uptime Institute's ongoing research into datacenter outage causes has determined that human error is a consistant root cause in many datacenter outages. Consequently, Uptime Institute created the Manangement and Operations (M&O) Stamp of Approval program, which helps organizations assess and improve datacenter management and operational behaviors while also ensuring that effective risk mitigation is in place.

IV. Excellence in Datacenter Management and Operations

Just as standards, requirements and certifications have evolved in many industries to maintain preparedness and ensure best practices, excellence in datacenter management and operations involves defining consistent and proven sets of standards, processes and measurements across an entire global datacenter portfolio.

As discussed, datacenter failures are almost never caused by one problem, and in most of the notable failures in recent decades, there was a breakdown or circumvention of established standards and certifications. It was not a lack of standards, but rather a lack of compliance or sloppiness, that most contributed to the disastrous outcomes. For example, at Delta, a glitch in the power supply was further exacerbated because some servers were not plugged into both the A and B sides of the power supply – demonstrating an inadequate or improperly followed installation process and poor oversight and/or risk-taking. This was then further complicated because recovery systems did not properly manage the reintroduction of services, so that databases became corrupted or untrusted. If leadership, operators and oversight agencies had adhered to their own policies and requirements and had not cut corners for economics or expediency, these disasters might have been avoided.

Human-related issues in running and operating datacenters typically involve one or more of the following:

1. Maturity of process – either too little or too much
2. Lack of process intergration
3. Reliance on people rather than process
4. Lack of coordination
5. Shift handover
6. Lack of process improvement
7. HR management
8. Inadequate levels of staff proficiency

Human-related issues in datacenter operations can be overcome with well-documented procedures, sufficiently mature and rigorously followed processes, staff traning and ongoing process improvement.

Ongoing operating and management practices and adherence to recognized standards and requirements, therefore, must be the focus of long-term risk mitigation. It is critical to pinpoint the elements that impact long-term datacenter performance, encompassing site management and operating behaviors, as well as documentation and mitigation of site-specific risks. It is also important to recognize that driving operational excellence across multiple datacenters is exponentially more difficult than managing just one. Technical complexity multiplies as you move to different sites, regions and countries where codes, cultures, climates and other factors are different.

Organizational complexity further complicates matters when the datacenters in a portfolio have different business requirements. With little difficulty, an organization can focus on staffing, maintenance planning and execution, training and operations for a single site. Managing a portfolio of datacenters turns the focus from projects to programs and from activity to outcomes. Processes become increasingly complex and critical. If you have more than one datacenter, it's about getting everyone on the same page and applying the same procedures and policies across the portfolio. The bigger the portfolio, the harder this is to do.

V. Benefits Accruing From Sustained Investment in Operational Excellence

It is virtually impossible for an organization's datacenter site culture, procedures and processes to be so refined that there are no details left unaddressed and no improvements that can be made. There is also a need to beware of hidden disparities between site policy and actual practice, as well as disparities between sites for multi-site operators.

Will a team be ready when something unexpected does go wrong? Just because an incident has not happened yet does not mean it will not ever happen. In fact, if a site has not experienced an issue, complacency can set in; steady state can get boring. Managers with foresight will create drills and get the team engaged with troubleshooting and implementing new, improved procedures. Furthermore, datacenter environments are never static, and the continuous review of performance metrics and vigilant attention to changing operating conditions is vital. Datacenter environments change frequently and if policies, procedures and practices are not revisited on a regular basis, they can quickly become obsolete. Just as 'good driver' discounts use an individual's track record as a reliable indicator of good ongoing behaviors (such as effective main-

tenance and safe driving habits), periodic datacenter recertification (biannually, at a minimum) provides a key indicator of ongoing effective facility management and operational best practices.

If adversity is the best teacher, then every failure in life is an opportunity to learn, and that certainly applies in the datacenter environment and other mission-critical settings. The value of undertaking failure analysis and applying lessons learned to continually develop and refine procedures is what makes an organization resilient and successful over the long term. There is tremendous value for organizations that hold themselves to a consistent set of standards over time, evaluating, fine-tuning and retraining on a routine basis. This discipline creates resiliency, ensuring that maintenance and operations procedures are appropriate and effective, and that teams are prepared to respond to contingencies, prevent errors and keep small issues from becoming large problems.

The bottom line is that effective datacenter operations reduce risk for both the MTDC providers and their customers. Well-informed datacenter management practices, systematic and documented policies and procedures, and effective training regimens, when applied and followed consistently and continually improved, can insulate a datacenter from the majority of downtime risks.

VI. Conclusions and Recommendations

Datacenters today are expected to be extremely reliable and provide maximum uptime for the IT and applications housed in them. However, such reliability involves much more than having a redundant datacenter infrastructure design. It also includes people, processes, operations and ongoing maintenance, as well as risk mitigation strategies. Datacenters are becoming increasingly complex to manage and operate, and many organizations where datacenters are not a core competency may not have the skillsets necessary for ongoing and effective management and operations of these facilities. On the other hand, operating datacenters *is* a core competency of MTDC providers. Their core business is in providing datacenter services, managing datacenter uptime and minimizing risks to customers. But an issue with establishing effective datacenter operations is that there really aren't comprehensive operations training facilities out there – only training for particular technologies. The result is that much of the training that needs to happen for datacenter operations is typically done in each datacenter, so it may not be consistent across multiple facilities.

So if you are going to leverage an MTDC provider, it is important to evaluate their processes and procedures, and ensure they are being followed, because this is a very difficult thing to manage across a large portfolio of facilities. MTDC providers should be able to walk customers and potential customers though practices for datacenter infrastructure management for each facility, including the upkeep of essential systems, particularly those supplying power and cooling, as well as their training programs for staff. Providers should have adequate monitoring and oversight of each facility and its important components to ensure that thresholds are being monitored and any incidents that crop up can be quickly mitigated before they cause an outage. To maintain the highest levels of availability, providers must also have adequate staffing and vendor support for each facility, as well as comprehensive, documented, and continually tested and improved policies and procedures. The bottom line is that no matter how resilient the design of a datacenter may be, outages are possible if the staff doesn't operate the facility with an eye toward operational excellence to maximize uptime.

CenturyLink®

Operational Excellence is One of Six Important Considerations When Choosing a Colocation Provider

The emergence of big data and the Internet of Things (IoT) is creating increasingly large quantities of data. In order to meet the needs of this growing data, businesses often have to decide how to securely deploy their IT infrastructure and store their mission-critical data for optimal density. Many on-premises datacenters are outgrowing their current capacity, dealing with outdated equipment and trying to keep up with shifts in technology, all while trying to expanding into other markets, both nationally and globally.

As technology continues to evolve and data requirements expand, businesses require additional bandwidth and increased network speeds just to keep pace. Without a solid foundation for their IT infrastructure, however, they risk costly outages and downtime that negatively affect their mission-critical applications.

COST CONSIDERATIONS

The most prominent options to combat this issue include expanding a current datacenter, building a new one or leasing datacenter space through colocation. Although businesses that build their own datacenter have complete control over their operating environment and have the capability to leverage existing space, doing so is often cost-prohibitive for smaller companies. Even businesses with enough capital to cover the up-front costs of building or expanding an on-premises datacenter don't always take into consideration the costs necessary for planning and design, property investment, power expenses, multi-level security, staffing and maintenance.

Colocation allows companies to focus on their primary business instead of investing extensive capital and resources in building and running a datacenter, and it continues to grow both in the United States and around the world. In fact, the datacenter colocation market is estimated to reach $33.2 billion by 2018 and grow to $54.13 billion by 2020.

PROVIDER PERKS

Colocation datacenters often serve as a bridge between on-premises datacenters and moving to the cloud. Whether they operate on-premises, in the cloud or in a hybrid IT environment, businesses can look to datacenters as an ideal disaster-recovery solution. They can utilize colocation to ensure business continuity within a realistic budget while creating a footprint in other markets or taking advantage of new business opportunities. Plus, consolidating operations and infrastructure into more efficient spaces reduces their capital and operating costs while achieving improved security and scalability.

Trusted and reliable datacenter providers enable their customers to enjoy 100 percent uptime, multi-level security, compliance and scalability with the flexibility to quickly react to business and market changes and easily expand or change services according to their business needs. They can enjoy the benefits a datacenter provides without the high costs and hassles of running one on-premises.

COMPLEX CHOICES

Choosing a colocation provider with the right datacenter experience and credentials to meet a company's unique needs can be a complex process. Although datacenters may not seem too different from one another, the way they are built and managed is essential to a business's mission-critical IT infrastructure and data. Utilizing the best processes and procedures for datacenter resiliency and achieving 100 percent uptime amid continuous IT change requires a culture of rigorous planning, testing and continuous improvement, something not all datacenter providers possess. One way of assessing a datacenter operator's day-to-day operations and datacenter management practices is to look for the Uptime Institute's Management and Operations Stamp of Approval. The M&O Stamp provides third-party validation that the operator is committed to operational excellence, providing its customers with confidence that they can expect an exceptionally high level of uptime and operational readiness from the datacenter operations team.

Companies need to ensure that the colocation provider they select has a history of operational excellence, secure environment options, system uptime to support backup and disaster-recovery plans and a state-of-the-art infrastructure that keeps their business applications and mission-critical data safe and ensures protection even if network operations are disrupted. Whether they need to be able to scale their infrastructure as their business changes or utilize remote support from on-site datacenter professionals, they should work with a provider that can meet their unique needs.

The geographic diversity available through top-tier colocation providers with a global footprint, including multiple locations within a metro or close proximity to one, ensures that businesses enjoy strong network connectivity between sites and helps maintain the reliability and uptime essential to business continuity and disaster recovery. If they need to get into a new market quickly, they can secure datacenter space with a colocation provider and have them migrate and install all the necessary equipment. Their IT infrastructure can be up and running, even in a market where they don't have a physical presence.

CenturyLink has over 20 years of experience in datacenter operations. Its operations management team averages more than 15 years of experience. Our deep experience and culture of continual improvement combine to make CenturyLink datacenters some of the best-run facilities in the world, as demonstrated by our 36 M&O Stamps of Approval covering all 53 CenturyLink-operated datacenters worldwide. Learn more about the six major areas to review when evaluating a colocation provider – operational practices, datacenter location, connectivity and service options, security and compliance, support services, and power and cooling – in our recently updated white paper, Six Important Considerations When Choosing a Colocation Provider. It is designed to help companies looking to invest in colocation ask the right questions to ensure an optimal deployment for their IT infrastructure. Businesses need to partner with a colocation provider that can meet their unique needs and scale with them as necessary. To learn more about the Six Important Considerations When Choosing a Colocation Provider, go to **www.centurylink.com/datacenterreport.**