Life Is On | Schneider Electric

Calculating ROI of Monitoring and Maintenance Service Contracts for New Data Centers

White Paper 152

Version 1

Energy Management Research Center

by Patrick Donovan and Wendy Torell

Executive summary

There are several strategies companies can deploy for maintaining their data center's physical infrastructure equipment (power, cooling) over its lifespan, including run-to-fail, run-to-alarm, calendar-based maintenance plans, and predictive or condition-based maintenance plans. While the value of a proactive maintenance plan provided by equipment OEMs or 3rd party service providers may be obvious to operations teams, quantifying this value in a credible way can be a challenge. In this paper, we provide a framework and introduce a tool for quantifying this value to help justify investment in monitoring and maintenance service contracts.

RATE THIS PAPER ★★★★★

Introduction

There are several different approaches to maintaining the physical infrastructure equipment (power, cooling, etc.) in a data center environment. White Paper 264, Attributes of an Effective Maintenance Program for Data Center Physical Infrastructure describes the various approaches to maintenance programs, and discusses the key attributes to look for in a service provider. It also describes how data analytics, digital services, and connected systems are enabling the evolution from calendar-based maintenance to condition-based maintenance.

In this paper, we provide a simple framework to help make the financial business case for the value provided by predictive or condition-based monitoring and maintenance contract. While service offers and their specific feature-sets vary from vendor to vendor, there are 3 core functions that generally apply in all cases:

- Remote 7x24 software monitoring of power and cooling infrastructure devices
- Predictive maintenance of power and cooling infrastructure devices
- Automated field service dispatch that includes replacement parts and "truck roll" costs.

Your actual return on investment (ROI) will be dependent on several factors, including the maturity of your IT operations, how critical the IT load is, how much time is spent on managing infrastructure, type of business, device failure rates, etc. These will determine if such a maintenance contract is a sound investment for a particular environment.

Next, we introduce a web-based ROI tool that applies this framework to prefabricated data center all-in-one modules. It helps you *quantify* this value and justify the investment. This tool offers a more credible prediction by giving the flexibility to adjust costs and other assumptions. Rather than prescribing the outcome, it provides a means for the user to experiment with "what-if" scenarios.

As Figure 1 illustrates, the framework includes 4 key categories of financial value: 1) a reduction in downtime occurrences (i.e., IT service outage), 2) a reduction in electricity consumption and the associated carbon emissions, 3) a reduction in staff and service costs related to power and cooling device management so their time can be re-allocated to other high priority tasks, and 4) the reduction of unpredictable spare parts costs throughout the data center's lifespan. In the sections that follow, we will dive into each of these categories.

Figure 1

A simple framework illustrating the value provided by monitoring & maintenance service contracts for data centers.



Downtime

Data center physical infrastructure systems tend to support critical IT loads, so most data center operators are averse to any downtime, especially unplanned outages. However, particularly when infrastructure equipment is newer and/or when staffing resources are constrained, the operation of the site might be done in a "run-to-alarm" fashion. In this mode, systems are left to operate with little to no proactive monitoring or preventative maintenance being performed. Operations teams only respond when there is a device alarm or noticeable change in equipment behavior.

This is a risky approach. Being unaware of device health, behavior, and environmental conditions will likely lead to more unexpected device failures and, worse, IT service outages. Being unprepared with little warning also means the duration of any service interruption will likely be longer than it would be otherwise. A lack of monitoring and regular maintenance also increases the risk of experiencing safety critical incidents.

Having remote visibility to power and cooling infrastructure devices, along with the environment the IT is housed in, is critical to reduce downtime and for rapidly diagnosing and overcoming service interruptions. Data center infrastructure management (DCIM) software, through its early warnings and device health assessments, can reduce the risk of unplanned service interruptions. In fact, according to an industry study, "Of the data centers that reported not having a DCIM system implemented, 80% also reported having experienced a major data center outage. By comparison, those sites that did have DCIM, roughly 29% reported having an outage."

However, effective use of DCIM requires training and on-going discipline by operations staff to use and maintain the software daily. If there are not enough staffing resources to do this, acquiring a service contract that includes "7x24" remote monitoring of your power and cooling infrastructure systems makes sense.

In assessing the potential financial risk of downtime both with and without a monitoring and maintenance service contract, there are four key variables that must be understood: (1) device failures, (2) device failures that lead to IT downtime, (3) the average time to restore IT service, and (4) the cost of IT downtime. By putting a value on each of these four key variables, you can estimate the impact of downtime with the following formula:

IMPACT OF DOWNTIME = NUMBER OF DEVICE FAILURES × % FAILURES CAUSING IT DOWNTIME × MEAN TIME TO RECOVER × DOWNTIME COST PER UNIT TIME

Each variable is explained below to provide guidance on estimating them for your business case.

Device failures

We define a device failure as a device state change that requires human intervention to ensure the device remains operational and fully capable within required parameters (i.e., battery runtime, heat rejection capacity). The percentage of the device population that fails each year is dependent on the device age and the operational environment, so it is important to understand the age mix of the fleet of infrastructure systems at the site. In general, the older the device population is and the harsher the operating conditions, the greater the failure rate. For example, the hotter the environment, the more battery failures will occur.

Device failures that lead to IT downtime

Not all device failures result in IT downtime. It's important to consider the percentage of power, cooling, and environmental monitoring infrastructure device failures that would result in an IT service outage. In some cases, this outage could be of an entire site, but other times, it could be a single rack. An example might be a UPS supplying power to a rack of IT and networking gear that shuts off due to a low battery condition causing the entire rack of IT to shut down. This could cause sales transactions at that site to come to a stop. Whether a failure leads to IT downtime or not depends largely on the following factors:

- criticality of the device to IT operations
- nature of the failure (e.g., does it cause the output power from the device to turn off or is it less urgent like a UPS fan failure?)
- level of power and/or cooling redundancy
- how long the device failure is allowed to persist (e.g., cooling unit shuts off allowing room to heat up over time, or a failed UPS battery self-test is ignored for too long), etc.

Average time to restore IT service

This is often referred to as mean time to recover (MTTR). This reflects how long the critical load is off-line, which has a major impact to the ultimate cost of the downtime event. Essentially, this is how long it takes your IT operations staff to detect the failure and implement corrective actions to restore IT operations again. Sometimes this duration is prolonged because the event failure isn't identified right away, the root cause it not figured out, and the solution to the failure is not resolved in a timely way. With a monitoring and maintenance service contract implemented, the source of the problem is likely identified and resolved faster compared to those operating in a "run-to-alarm" fashion. The above-mentioned industry study also found "the presence of DCIM reduced the average length of an outage from 342 minutes without DCIM compared to 36 minutes with DCIM". That represents a 90% reduction in outage duration.

Cost of IT downtime

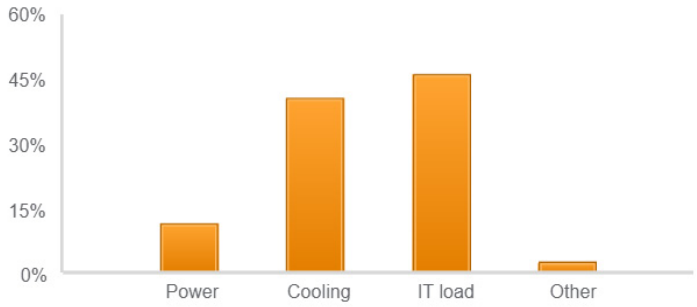
The cost of downtime per minute can vary widely and is highly dependent on your specific case. It varies, in part on your industry, and the type and amount of IT service(s) and business functions impacted by the failure (e.g., data storage vs. ecommerce services vs. operations at a distribution center). This eBook by Carbonite explores the impact of downtime and data loss on enterprise IT departments and provides some data on the costs.

Electricity & Emissions

There are many factors that affect the electricity consumption of a data center, including power and cooling system design choices (e.g., the UPS efficiency, the cooling architecture, system redundancy) and operational practices (e.g., operating set points like temperature, hours of economizer mode operation, load on the systems, maintenance practices). Figure 2 is a typical breakdown of a data center's energy.

Figure 2

A typical breakdown of energy in a data center. Note, this varies by data center architecture and design choices. See TradeOff Tool, Data Center Efficiency & PUE Calculator.



Since the cooling system often represents the biggest portion of the total energy consumption in a typical data center next to IT load energy, and because mechanical systems are more likely than electrical systems to become inefficient over time without maintenance, this framework quantifies the monitoring/maintenance energy cost and emissions reduction benefit based only on the cooling system. Three examples of electrical inefficiencies that occur over time with cooling systems are below:

- Cooling system filters and heat rejection units can become dirty and clogged which reduces airflow and forces an increase in energy consumption in order to reject the same amount of heat.
- Grease in bearings can dry or leak out which increases friction and leads to higher electric losses.
- Outdoor condensers can become covered in debris, dust, tree pollen, leaves, insects, etc. Dirty condenser coils result in higher system pressure and higher current draw, ultimately leading to greater energy expense.

The energy cost impact of a data center with no maintenance can be calculated using the following formula, where the **cooling energy factor** includes the energy penalty for years with no maintenance contract.

$$\text{ENERGY EXPENSE} = \left[\text{COOLING BASELINE ENERGY} \times \text{COOLING ENERGY FACTOR} + \text{POWER SYSTEM ENERGY} + \text{OTHER SYSTEM ENERGY}\right] \times \text{ELECTRICITY COST PER KWH}$$

Likewise, the carbon emissions impact can be calculated using a similar formula but multiplying the total energy by the electricity emission factor.

$$\text{ENERGY EMISSIONS} = \left[\text{COOLING BASELINE ENERGY} \times \text{COOLING ENERGY FACTOR} + \text{POWER SYSTEM ENERGY} + \text{OTHER SYSTEM ENERGY}\right] \times \text{CO}_2\text{ EMISSION FACTOR (KG PER KWH)}$$

Cooling energy factor

A study by Oak Ridge National Laboratory operated by the US Department of Energy[1], measured the impact of air-filter condition on HVAC equipment. Although the study is not recent, the physics leading to the quantified energy penalties measured and reported in the study have not changed. The study states that "as the air filter becomes clogged, the airflow across the evaporator is reduced, and heat transfer from the air to the evaporator is inhibited. A dirty evaporator surface also lowers the

[1] Oak Ridge National Laboratory, Impact of Air-Filter Condition on HVAC Equipment, L. Jung, 1987

heat-transfer coefficient so that less heat can be transferred to the refrigerant. When either the filter becomes clogged or the evaporator surface becomes dirty, the evaporator temperature and the suction pressure are lowered, thus reducing the [refrigeration] cycle efficiency."

The study calculates the energy factor for nine cases where the systems are performing in a sub-optimal state (e.g., very dirty filter, airflow blockage, dirty evaporator, dirty condenser). Table 1 provides recommended energy penalties to apply to the cooling system energy in years 2 through 5 of a data center's operation when no maintenance contract is in place. Since the data center environment and operational procedures can greatly impact what state a cooling system would be in, year-by-year, the values are only a guide, and should be adjusted accordingly to reflect your particular environment, i.e., whether it be considered highly clean or dirty.

Table 1

Energy penalties for years 2 through 5 of a data center's operation without proactive maintenance, derived from Oak Ridge National Laboratory study, Table 2, interpolating case 4 for year 2 and case 7 for year 5.

Year	Capacity factor	Power factor	Cooling energy factor @ 100% load
1	1	1	100% (baseline)
2	0.93	1.11	119%
3	0.88	1.13	129%
4	0.83	1.16	139%
5	0.78	1.18	151%

** Note, Cooling energy factor is calculated as Power factor divided by Capacity factor; Capacity factor increases at loads less than 100%, which decreases the energy factor, so operating load should be considered.*

Electricity cost

The electricity cost per kWh in your data center's geography should be used to calculate the cost of the energy consumed. Global Petrol Prices is a helpful source for finding national average electricity prices. This can be used as a starting point in calculating the electricity cost savings, but if you have specific rates from your utility provider, that will provide a more accurate reflection of the potential electricity cost savings.

CO2 emission factor

The CO2 emission factor is the amount of greenhouse gases (GHG) emitted in kilograms per kilowatt hour of electricity consumed. As a starting point, Carbon Footprint, LTD has data that can be downloaded on the average emission factors by country. This data is continually updated, as utility providers are always evolving their mix of energy sources. You can also go to your utility provider's website to get their mix of energy sources and calculate your electricity emission factor using our Electricity Emission Factor Calculator, for a more accurate reflection of the potential carbon emissions savings.

Staffing

Managing data center infrastructure generally relies on two types of staffing to ensure systems are operating and maintained appropriately. This includes:

- IT operations staff – These are the people responsible for the day-to-day operations of the IT equipment and the supporting infrastructure across the site.
- Field service staff – These are the technicians who respond when equipment failures occur, or preventive maintenance (e.g., battery replacement) is needed. They may be employees of the data center, the equipment vendor, or be a third-party service provider (e.g., managed service provider, or MSP).

Without remote software visibility, it takes IT operations staff longer to find, diagnose, and resolve problems. Those that don't proactively monitor or have a vendor do it, and instead rely on a "run to alarm" strategy, will spend significant time responding to issues as they come up unexpectedly, trying to find the root cause, calling vendor support, ordering parts, and performing swap outs.

Without data from DCIM and the expertise of experienced vendor support teams, there is an increased likelihood of unnecessary visits by field service staff. Examples of common situations include:

- Device LED is flashing or making an audible alarm, and people onsite assume there's a problem and call for service.
- UPS shuts down due to operating on battery for too long and staff just ask for it to be replaced even though the UPS may have been performing as it was supposed to.
- IT loads shut down due to a simple overloaded circuit but without DCIM software to provide visibility as to why, a technician service call is initiated.
- Cleaning crew plugs in a vacuum after hours tripping a breaker that causes the UPS to go to battery and eventually shut down the IT load.

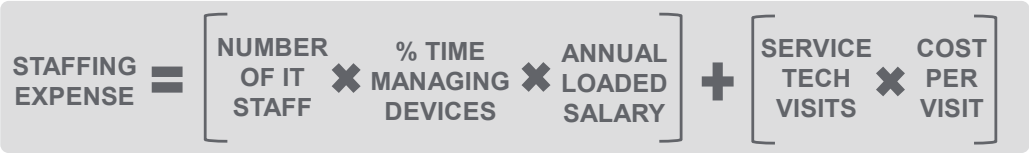
A monitoring and maintenance contract improves staff efficiency by putting the task of constantly monitoring the health and status of your power and cooling systems in the hands of a third party. Outsourcing this function to experts makes you more effective and proactive in avoiding problems that might, otherwise, lead to an IT service interruption. And when problems do occur, reacting to incidents, determining a root cause, and resolving the issue will all be faster.

The maintenance portion of the contract typically means prompt response by service personnel who arrive on site with the correct type and quantity of needed parts. More effective service contracts will even automate field service dispatch. And vendors who use more effective DCIM solutions will be more likely to detect changes and aberrations in device behavior and operating conditions that eventually lead to service interruptions if it went unnoticed. Thus, an effective monitoring and maintenance service contract allows your IT staff to focus more on strategic tasks, be less distracted, and will reduce the frequency of service visits that can be disruptive to your IT staff.

For the purposes of calculating an ROI for monitoring and maintenance service contracts, the framework described in this paper quantifies the value of this staff efficiency gain in terms of financial savings. However, this is not meant to imply that labor efficiency should equate to a reduction in staff count.

There are three costs that should be understood about your staff operations when assessing the financial impact of staff efficiency improvements by implementing

monitoring and maintenance service contract. You should first estimate these staffing costs when operating in a "run to alarm" mode (i.e., no DCIM, remote monitoring, or maintenance contract) to serve as a baseline. Depending on the maturity of your operations and maintenance program, the service contract will have a varying improvement in each of the three categories. Think about past incidents and experiences, and what percent of them could have been avoided or the time spent dealing with a failure could've been avoided with the visibility and insights provided by third-party remote monitoring and the speed of automated field service dispatch. The formula to estimate the staffing expense in managing the power and cooling infrastructure is as follows:



Each cost is explained below to provide guidance on estimating them for your business case.

Operations IT staff cost

The operations IT staff expense is based on (1) the number of IT staff needed to support the entire data center; (2) the average total annual cost of each IT operations staff employee including insurance and other related benefits; and (3) the percentage of the overall operations IT staff time that they spend managing and maintaining these devices. This includes time spent troubleshooting alarms, resolving calls and emails about related problems, managing device consumables (e.g., UPS batteries), ordering parts, deploying new or replacing old units, and other asset management tasks.

Note, in the United States, IT operations technicians have a typical base salary ranging from $53k - $70k[2]. Benefits are valued at approximately 30% of total compensation[3], so this puts a typical total compensation for IT operations staff at $100k per full time person.

Field service staff cost

It is important to estimate the number of service visits (sometimes referred to as "truck rolls"), made across all sites to address issues with physical infrastructure devices. This should include visits made for both actual device failures, as well as for issues where no real problem was found. Service visits are often made when no real problem exists or is solvable through simple adjustments in unit operating parameters.

An effective monitoring and maintenance service contract can prevent many of these unnecessary service visits. For each visit, you need to estimate the average cost to have a field service technician (either yours or from an OEM or 3rd party) arrive on site and diagnose the problem. The cost should include the technician wages and travel costs. According to the Technology Service Industry Association, truck rolls cost the average business around $1,000 per dispatch[4]. This value can vary depending on the type and size of devices being serviced, and the distance the technician(s) must travel.

[2] https://www.salary.com/research/salary/recruiting/it-operations-technician-salary

[3] https://www.bls.gov/news.release/ecec.nr0.htm

[4] https://helplightning.com/blog/reduce-truck-roll-costs/

Spare Parts

Another financial savings to include in the ROI of maintenance contracts is spare parts cost reduction. Without a maintenance contract in place, the cost for needed spare parts due to failures can be an unpredictable expense since it is difficult to plan for when a failure will occur or when consumable parts like batteries reach their end of life. Many maintenance service contracts include the cost of these parts or whole system replacements, thus avoiding the unanticipated spikes in opex when they become necessary.

Using the number of device failures (mentioned in the **Downtime** section) each year and the average parts cost anticipated per failure event, the spare parts cost reduction can be calculated. Note, when spare parts are no longer available for a particular subsystem, this is generally an indicator it may be time to modernize or replace the system, to avoid future downtime.

A value framework tool

Using the framework above, and applying it to prefabricated all-in-one data center modules up to 200 kW in capacity, we created a freely available, web-based TradeOff Tool (**Figure 3**), Prefabricated Data Center Service Contract ROI Calculator, for quickly estimating and understanding the potential financial value of a monitoring and maintenance service contract based on your specific inputs. Note, although the tool is specific to prefabricated modules, the logic and framework applies more broadly. The tool was scoped for simplification, since a tool where many more combinations of power and cooling architectures, capacities, redundancies, etc. are allowed makes estimating service contract cost more complex. We propose using the tool as a starting point for larger or non-modular data center projects, and scaling the cost based on capacity to give you a rough directional ROI.

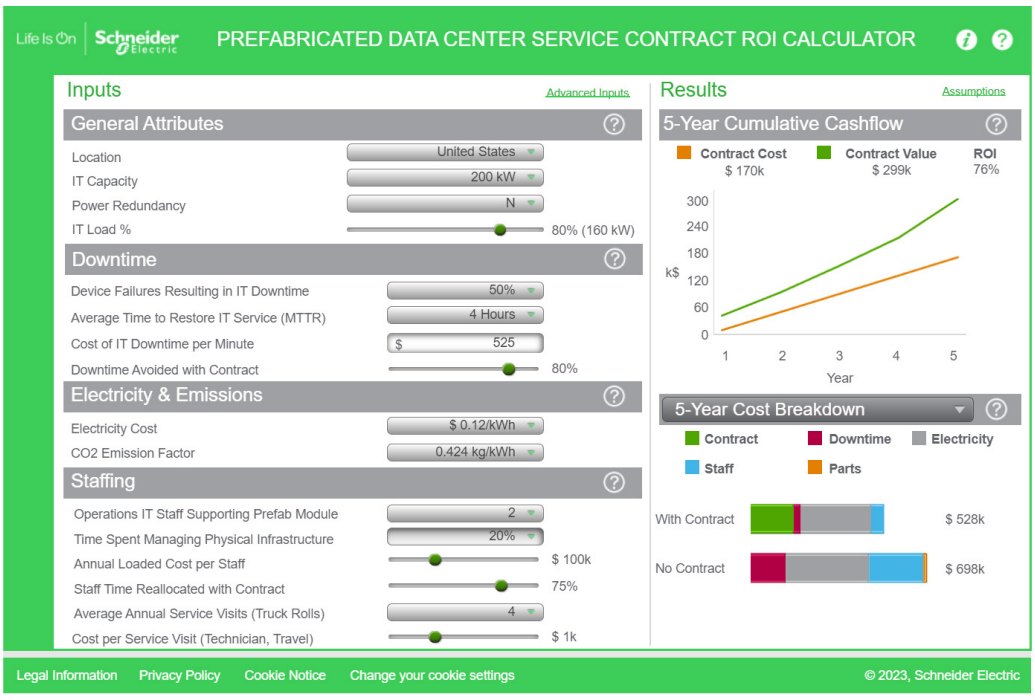


Figure 3

TradeOff Tool to provide a framework and quantify the value of Monitoring & Maintenance Contracts for Prefabricated Data Centers.

Methodology

In this TradeOff Tool, the quantified value of a monitoring and maintenance service contract is derived from the four categories presented earlier in this paper: 1) a reduction in downtime (i.e., IT service outage) occurrences, 2) a reduction in electricity cost and CO2 emissions, 3) staff optimization, and 4) avoidance of unanticipated parts cost.

The calculated outputs are based on user-selected inputs about their data center, including the capacity, IT load, redundancy, cost of downtime, cost of electricity, cost of staff, and so on. Default values are assumed to be reasonable values based largely on customer anecdotes and location-specific data. Since actual values depend significantly on the data center's attributes and environment, it is important for the user to choose inputs that reflect actual operational experience and site conditions.

The "Advanced Inputs" tab gives users the ability to further refine the results by adjusting additional key variables (device failure rates, cooling energy penalties without maintenance, and annual contract costs) based on their specific expectations and/or needs. Variables not adjustable are noted in the "assumptions" link.

Results

The tool calculates a 5-year cash flow and shows an ROI percentage. The chart illustrates the cumulative cashflows of the contract cost vs. the value gained from the contract (see Figure 3). Cumulative means that each year's amount reflects the previous years' amounts plus the amount in that year. These amounts are adjusted to reflect a net present value (NPV) using the selected cost of capital.

The orange line represents the cost of the monitoring/maintenance contract over the 5 years. The green line represents the value that the service is expected to provide based on the inputs and assumptions used in the model. It is a summation of the downtime events avoided, electricity cost savings, a reduction in staff time spent diagnosing and maintaining power and cooling devices, and the value of eliminating spare parts cost.

A 5-year cost breakdown is also provided to demonstrate the cumulative downtime costs, energy costs, staff costs, and parts cost with and without the contract implemented. Having this side-by-side comparison of each category helps illustrate where the service has the most impact for the scenario entered.

We ran three scenarios of a 200 kW data center with 1N redundancy, to illustrate how the value of the service contract can vary (from 11% ROI to 180% ROI), and what some of those key drivers are. Figure 4 presents these scenarios and shows how cost and duration of downtime, energy penalties, and cost of staffing/servicing can affect the ROI or payback. See Appendix for further details of each scenario.

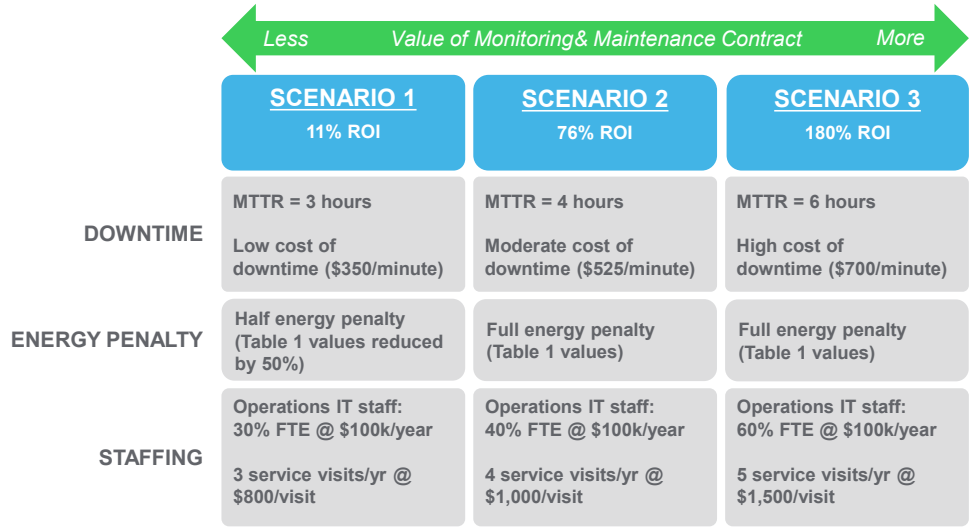


Figure 4

Three scenarios showing how monitoring & maintenance service contract value varies based on downtime, energy penalties, and staffing baseline conditions.

See Appendix for screenshots of TradeOff Tool showing details of each scenario.

Conclusion

Consistent monitoring and proactive maintenance of data center physical infrastructure systems is necessary to ensure uninterrupted IT operations. However, doing this well takes discipline, time, expertise, and good data. Particularly for resource-constrained IT operations teams, many end up operating in a “run to alarm” mode where action and attention is only given when a problem occurs. Operating in this way risks IT service interruptions and likely means longer MTTR periods than it otherwise would be if the data center was well managed and maintained.

Data center managers and operators should consider monitoring and maintenance service contracts by qualified third party vendors. These contracts typically provide three core functions:

- Remote 7x24 software monitoring of power & cooling infrastructure devices
- Predictive maintenance of power and cooling infrastructure devices
- Automated field service dispatch that includes replacement parts and “truck roll” costs.

Well executed by infrastructure experts with effective digital tools, these services can reduce downtime, reduce energy consumption, increase staff efficiency, and reduce or eliminate replacement part costs. But justifying the investment in a service contract with an ROI can be difficult to do in a credible way. In this paper, we describe a framework for calculating the return on investment compared to operating in a “run to alarm” mode.

Data center managers and their IT operations teams can use the newly developed TradeOff Tool to quantify the estimated value of implementing a monitoring and maintenance service contract. Although the tool today is aimed at smaller (200 kW and less) prefabricated, all-in-one data center modules, the framework applies for any data center based on their individual cost of downtime, cost of service, projected number of incidents, and other user-specified attributes.

About the authors

Patrick Donovan is a Senior Research Analyst for the Energy Management Research Center at Schneider Electric. He has over 28 years of experience developing and supporting critical power and cooling systems for Schneider Electric’s Secure Power Business unit including several award-winning power protection, efficiency, and availability solutions. An author of numerous white papers, industry articles, and technology assessments, Patrick's research on data center physical infrastructure technologies and markets offers guidance and advice on best practices for planning, designing, and operation of data center facilities.

Wendy Torell is a Senior Research Analyst at Schneider Electric’s Energy Management Research Center. In this role, she researches best practices in data center and building design and operation, publishes white papers & articles, and develops TradeOff Tools to help clients optimize the availability, efficiency, and capex/opex costs of their facilities. She also consults with clients on availability science approaches and design practices to help them meet their performance objectives. She received her bachelor’s degree in Mechanical Engineering from Union College in Schenectady, NY and her MBA from University of Rhode Island. Wendy is an ASQ Certified Reliability Engineer.

RATE THIS PAPER ★★★★★

Resources

Attributes of an Effective Maintenance Program for Data Center Physical Infrastructure
White Paper 264

Benefits of shifting from traditional to condition-based maintenance in electrical distribution equipment
White Paper

Browse all white papers
whitepapers.apc.com

Prefabricated Data Center Service Contract ROI Calculator
TradeOff Tool 36

Browse all TradeOff Tools™
tools.apc.com

Note: Internet links can become obsolete over time. The referenced links were available at the time this paper was written but may no longer be available now.

Contact us

For feedback and comments about the content of this white paper:

Schneider Electric Energy Management Research Center
dcsc@schneider-electric.com

If you are a customer and have questions specific to your data center project:

Contact your Schneider Electric representative at
www.apc.com/support/contact/index.cfm

© 2023 Schneider Electric. All rights reserved.

Appendix

Three sets of screenshots of the TradeOff Tool are shown below to provide details on the three scenarios presenting varying ROI, as summarized earlier in **Figure 4**.

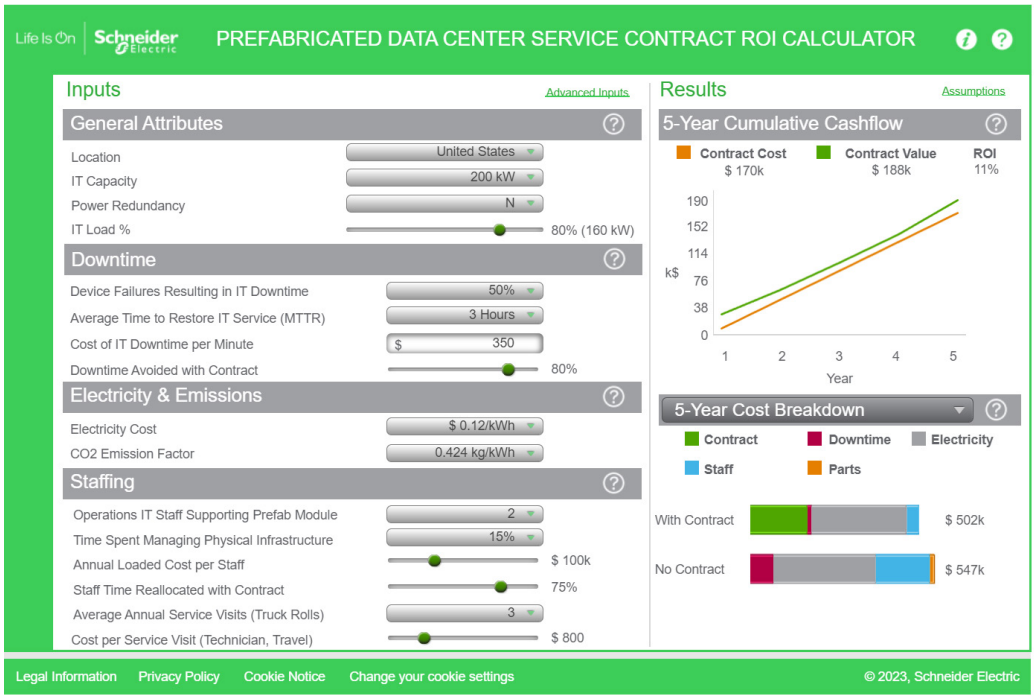


Figure A1a

Scenario 1, showing ROI of 11%

Main inputs

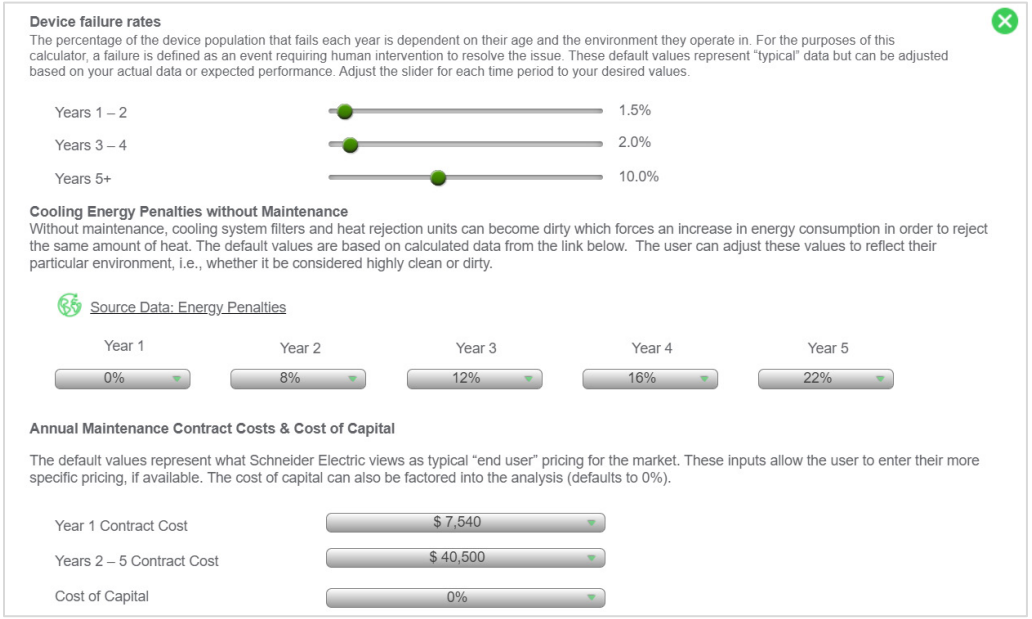


Figure A1b

Scenario 1, showing ROI of 11%

Advanced inputs

Figure A2a

Scenario 2, showing ROI of 76%

Main inputs

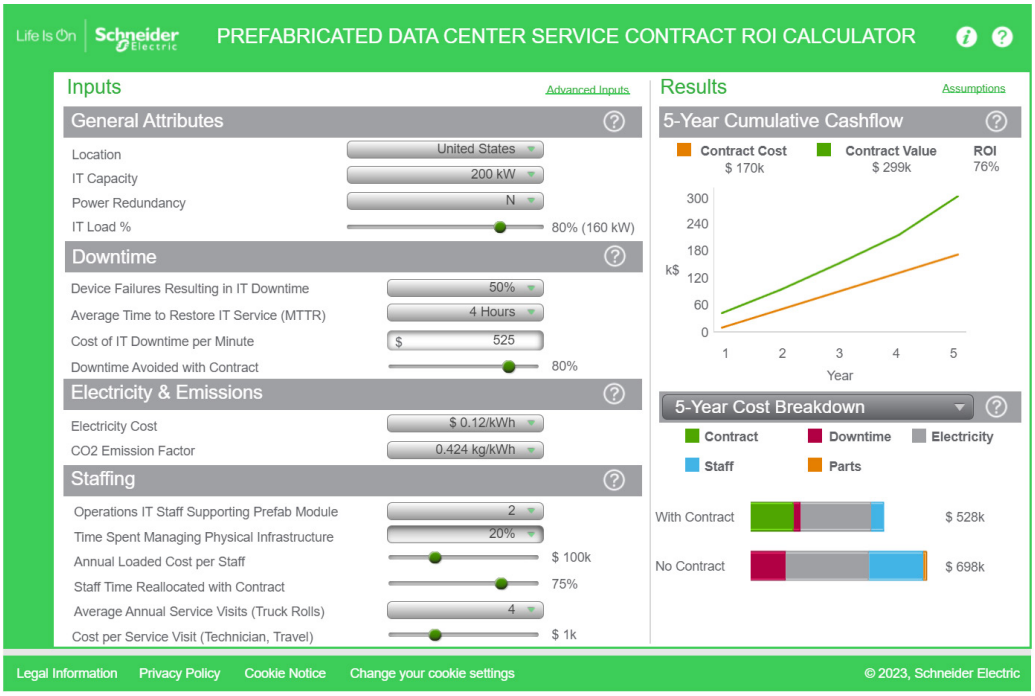
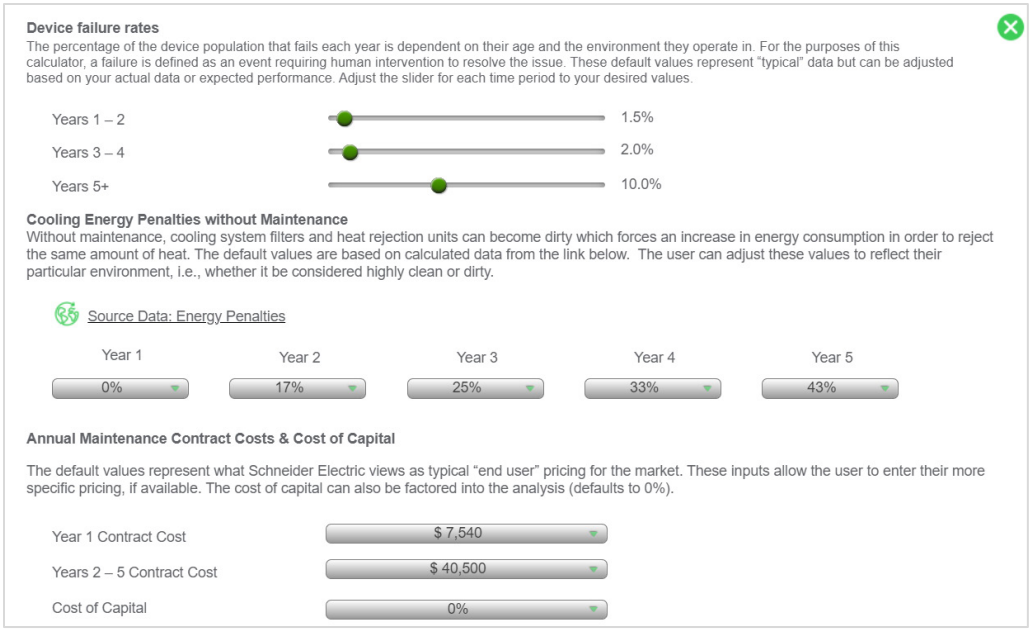


Figure A2b

Scenario 2, showing ROI of 76%

Advanced inputs



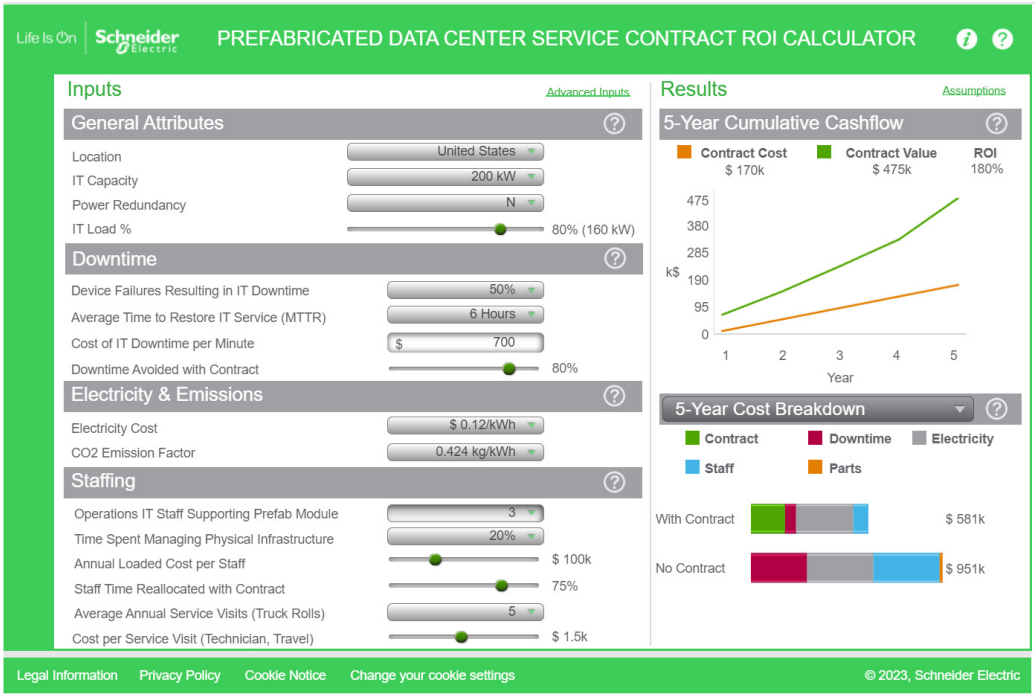


Figure A3a

Scenario 3, showing ROI of 180%

Main inputs

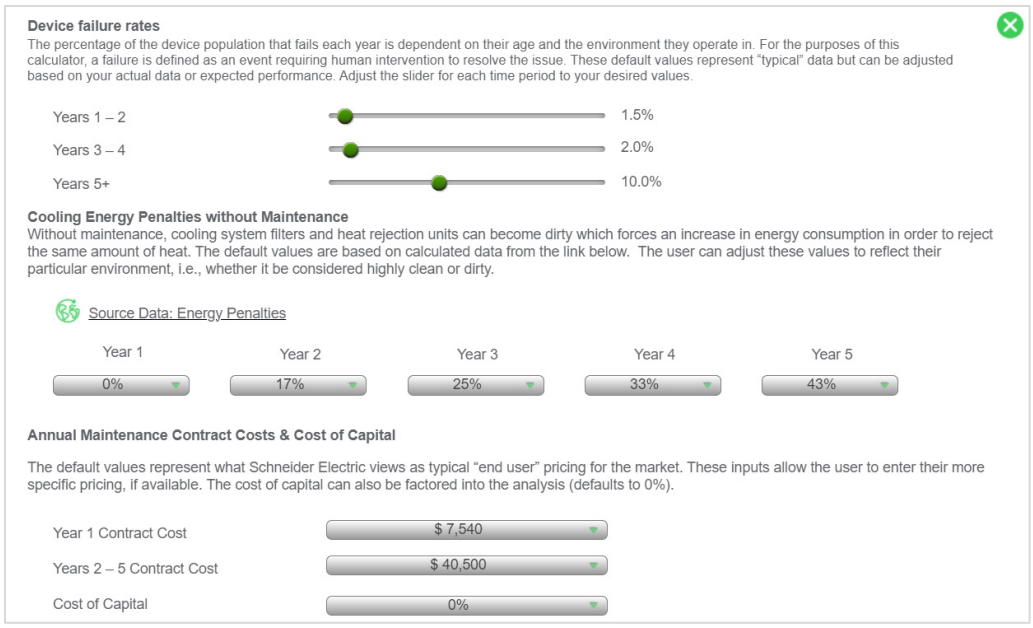


Figure A3b

Scenario 3, showing ROI of 180%

Advanced inputs