

Calculating ROI of DCIM Monitoring for Distributed IT Sites

White Paper 292

Version 1

Energy Management Research Center

by Patrick Donovan and Wendy Torell

Executive summary

Data Center Infrastructure Management (DCIM) software is widely used to improve resiliency, security, and staff efficiency. The urgent need to remotely monitor critical power and cooling infrastructure with DCIM is, perhaps, obvious to IT Operations teams tasked with maintaining availability of sprawling IT portfolios. However, calculating an ROI or payback analysis for executives can be difficult to do credibly for software management tools. In this paper, we provide a framework and introduce a tool for quantifying this value to help justify investment in DCIM monitoring and alarming functions.

RATE THIS PAPER



Introduction

Data center environments have become increasingly complex, with information technology (IT) spanning on-premise, centralized data centers, local and regional edge sites, as well as systems, application, and data being hosted by colocation and cloud service providers. Managing these distributed assets is a challenge, especially in sites without IT staff onsite. In White Paper 281, [How Modern DCIM Addresses CIO Management Challenges within Distributed Hybrid IT Environments](#), we discuss the growing challenges faced by CIOs today and how Data Center Infrastructure Management (DCIM) has evolved to help address these challenges.

DCIM software tools offer many features and benefits. Specifically, the monitoring & alarming functions provide visibility of all monitored devices, through the DCIM software interface, to have awareness of device status, health, and their environmental conditions. Based on user-defined thresholds and settings, notifications and alarms are reported in the DCIM software. The software provides additional means to notify operations staff by email or texts. This DCIM monitoring software is typically accessible through a mobile device app or a web browser. The monitoring & alarming functions also provide staff with the ability to remotely manage their fleet of devices (i.e., update firmware, adjust alarm thresholds and network security settings, etc.).

Even though software management tools like DCIM often intuitively provide value to users, sometimes justifying the purchase with a financial return on investment (ROI) or payback analysis can be difficult to do in a *credible* way. This is, in part, because calculating the value of software requires estimating future outcomes on how your operations would change once the software is implemented.


In this paper, we provide a simple framework to help make the business case for the value provided by DCIM monitoring and alarming functions. Your actual values will be dependent on several factors, including the maturity of your IT operations, how critical the IT load is, type of staffing present at each site, type of business, type and age of IT, and supporting infrastructure implemented, etc. These will determine if DCIM is a sound investment for a particular environment.

Next, we introduce a web-based tool that applies this framework to help 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 3 key categories of financial value: 1) a reduction in downtime occurrences (i.e., IT service outage), 2) a reduction in staff and service costs related to power and cooling device management, and 3) a reduction in security and environmental incidents. In the following sections, we will dive into each of these categories.

Figure 1

A simple framework illustrating the value provided by DCIM monitoring & alarming functions in a distributed IT environment.








Downtime

Data center physical infrastructure systems (power, cooling, distribution, etc.) support critical IT loads, so most data center operators are averse to any downtime. Particularly with portfolios of highly distributed, unmanned edge computing sites, maintaining resiliency can be a challenge (see **Figure 2**).

Figure 2

Distributed, hybrid IT environment makes it more challenging to maintain resiliency



-  Having visibility to many geographically dispersed sites
-  Lacking onsite staff
-  Managing alarm storms
-  Maintaining a large fleet

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. Being unaware of device health and environmental conditions will likely lead to more unexpected device failures and, worse, IT service outages. DCIM, 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.”

This impact of DCIM on downtime will vary based on how mature your operations and maintenance teams are. This includes:

- staffing coverage
- level of staff training & knowledge
- level of DCIM use

In assessing the potential financial risk of downtime both with and without DCIM, 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:

$$\text{IMPACT OF DOWNTIME} = \text{NUMBER OF DEVICE FAILURES} \times \% \text{ FAILURES CAUSING IT DOWNTIME} \times \text{MEAN TIME TO RECOVER} \times \text{DOWNTIME COST PER UNIT TIME}$$

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 across the hybrid environment. In general, the older the device population is and the harsher the operational conditions, the greater the failure rate.

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 is not figured out, and the solution to the failure is not resolved in a timely way. With DCIM implemented, the source of the problem is often identified and resolved faster. 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. e-commerce services vs. operations at a distribution center).

Staffing

Managing distributed IT infrastructure generally relies on various 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 sites. They are often centrally located at a distance from distributed sites.
- **On-premise (non-IT) staff** – These are people with core functions not related to supporting the IT and physical infrastructure equipment (e.g. warehouse personnel, retail store salespeople, branch managers), who sometimes are compelled to respond to “urgent” or “unknown” situations that occur with the equipment because they are the only ones on site.
- **Field service staff** – These are the technicians who respond when equipment failures occur, or preventive maintenance (e.g., battery replacement) is needed.

Manual monitoring without DCIM software can be very time consuming for the **IT operations staff**. Without remote software visibility, it takes longer to find, diagnose, and resolve problems. Unnecessary time is spent sending IT staff from site to site for things like misdiagnoses and firmware updates. Those that don’t proactively monitor, and instead rely on a “run to fail” strategy, also spend significant time responding to issues as they come up unexpectedly and performing swap outs.

Blinking LEDs, beeping alarms, loss of service due to a device shutdown, etc. are all things that trigger support calls that take **non-IT staff** attention away from what their primary job function is. These interruptions result in a loss of productivity that costs the organization money.

There is also an increased likelihood of unnecessary **service technician visits**. 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.

Implementing DCIM improves staff efficiency by providing your IT operations team with remote visibility into device health and environmental conditions at your distributed sites. This visibility makes you more effective at proactively managing your assets to avoid problems and allows you to be more effective at preventing, reacting to and resolving incidents. **Thus, it allows your IT staff to focus more on strategic tasks, reduces distractions preventing your non-IT staff from their normal job functions, and reduces the frequency of service visits.**

For the purposes of calculating an ROI or payback, this framework does quantify the value of this 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

DCIM. You should first estimate these staffing costs when no DCIM is present, to serve as a baseline. Based on the maturity of your operations and maintenance program, DCIM 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 with the visibility and insights provided by DCIM. The formula to estimate the staffing expense in managing the power and cooling infrastructure is as follows:

$$\text{STAFFING EXPENSE} = \left[\text{NUMBER OF IT STAFF} \times \% \text{ TIME MANAGING DEVICES} \times \text{ANNUAL LOADED SALARY} \right] + \left[\text{NON-IT STAFF HOURS SPENT} \times \text{OPPTY COST PER HOUR} \right] + \left[\text{SERVICE TECH VISITS} \times \text{COST PER VISIT} \right]$$

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

IT operations staff cost

The IT operations staff expense is based on (1) the number of IT operations staff needed to support the entire distributed enterprise portfolio; (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 IT operations 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), 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 - \$70k¹. Benefits are valued at approximately 30% of total compensation², so this puts a typical total compensation for IT operations staff at \$100k per full time person.

On-premise (Non-IT) staff cost

The non-IT staff expense is based on the total number of hours spent per year by non-IT staff having to deal with power and cooling infrastructure devices. Examples include personnel who notice something wrong or unusual about the back-office equipment. You need to then quantify the value of this time spent. Note, that value should reflect the opportunity cost of the function they *should* be doing instead of reacting to the problem-at-hand. This is generally significantly more than the salary of the individual(s). For example, the value of a retail salesperson is not just the hourly wage, but also the value of the lost sales orders. This metric is sometimes referred to as “sales per labor hour” or “SPLH” and can vary greatly by industry.

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. Particularly since distributed IT sites are often unmanned or do not have technical people on site, service visits are often made when no real problem exists or is solvable through simple adjustments in unit operating parameters.

¹ <https://www.salary.com/research/salary/recruiting/it-operations-technician-salary>

² <https://www.bls.gov/news.release/ecec.nr0.htm>

Remote monitoring and management through DCIM software 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 a 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 dispatch³. This value can vary depending on the type and size of devices being serviced, and the distance the technician(s) must travel.

Security & environmental

Modern DCIM software monitoring tools integrate with environmental monitoring (temperature, humidity, and leak sensors), as well as security cameras that monitor room and rack access for unauthorized access to the IT. Some, like Schneider Electric, also offer a cybersecurity assessment feature that will perform regular scans of connected devices to alert you of out-of-date device firmware and vulnerable network security settings that might increase the risk of a successful cyberattack. These DCIM tools reduce the risk of:

- Cybersecurity breaches
- Physical security breaches
- Environmental incidents

The cost of a single security or environmental incident can exceed the cost of the DCIM solution by orders of magnitude, making it overwhelmingly dominate the value calculation. For example, while a cyber breach for a small business has been reported to cost \$108k USD⁴, the global average cost for a data breach has been reported to be as high as \$4.35 million⁵. There are many factors affecting what this value would be for a particular site, such as how large the breach is and what data is compromised. As another example, an unnoticed water leak can cause enough damage to quickly exceed the cost of DCIM, more than justifying the entire expense of the solution.

While security and environmental damage generally dominates the calculation, for “what-if” planning, some companies like to include their value, as a way to point out to stakeholders the potential value of the software, as some are not aware of DCIM’s cybersecurity assessment or environmental monitoring functionality⁶. For those that choose to include it, we recommend using your estimated value for a single incident.

A value framework tool

Using the framework above, we created a freely available, web-based TradeOff Tool (**Figure 3**), [DCIM Monitoring Value Calculator for Distributed IT](#), for estimating and understanding the potential value of DCIM monitoring tools based on your specific inputs. Users can quickly quantify the estimated value of the software tools’ monitoring & alarming functions for distributed, hybrid IT environments. Note, the value of DCIM planning and modeling functions (i.e., mapping of asset dependencies, simulating adds/moves/changes, energy efficiency modules, IT server optimization, [CFD](#) analysis, etc.) are not accounted for in this tool.

³ <https://helplightning.com/blog/reduce-truck-roll-costs/>

⁴ <https://prowritersins.com/cyber-insurance-blog/average-cost-of-a-data-breach-for-small-businesses/>

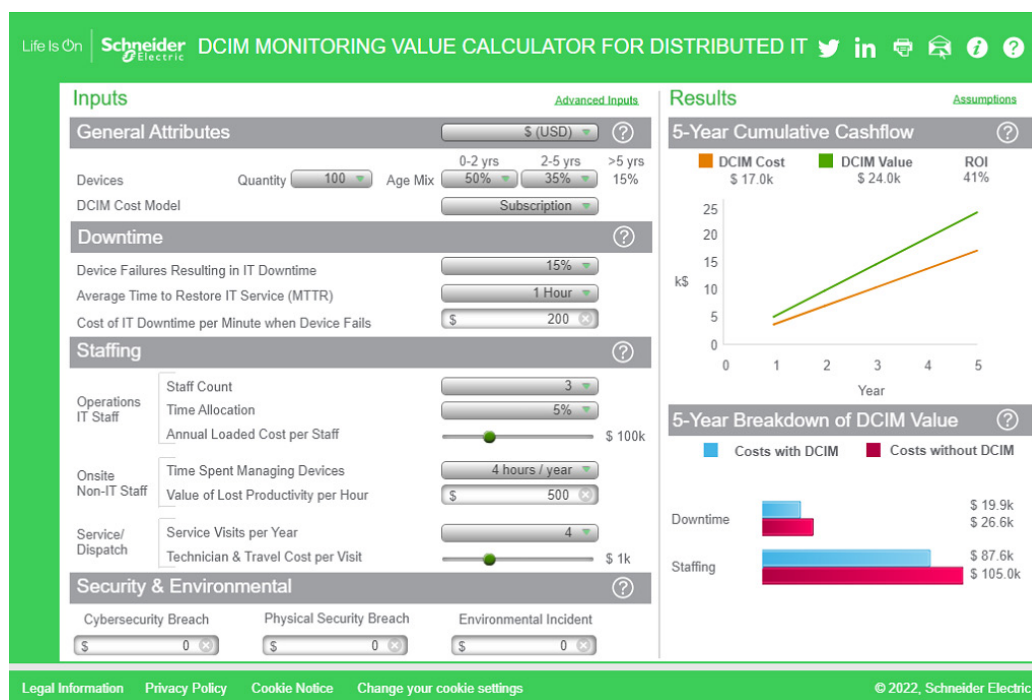
⁵ <https://www.securitymagazine.com/articles/98486-435-million-the-average-cost-of-a-data-breach>

⁶ White Paper 216, [Cybersecurity Guidance for Data Center Power and Cooling Infrastructure Systems](#)

Figure 3

TradeOff Tool to provide a framework and quantify the value of DCIM monitoring

<https://www.se.com/ww/en/worksolutions/systems1/data-center-and-network-systems/trade-off-tools/dcim-monitoring-value-calculator-for-distributed-it/>



Methodology

In this TradeOff Tool, the quantified value of DCIM monitoring & alarming is derived from the three categories presented earlier in this paper: 1) a reduction in downtime (i.e., IT service outage) occurrences, 2) a reduction in staff and service costs related to power and cooling device management, and 3) a reduction in security and environmental incidents.

The calculated outputs are based on user-selected inputs about their sites, number of monitored devices, staff costs, cost of downtime, DCIM solution costs, and so on. Default values are assumed to be reasonable values based largely on customer anecdotes. Since actual values depend significantly on the site's attributes and maturity, it is important for the user to choose inputs that reflect their actual operational experience and site conditions. The "Advanced Inputs" tab gives users the ability to further refine the results by adjusting additional key variables based on their specific expectations and/or needs. Variables not adjustable are noted in the "assumptions" link.

Results

The tool presents results based on your selection of preferred cost model: Perpetual license-based DCIM or subscription-based DCIM. Perpetual license involves paying a one-time upfront cost to use the DCIM software in perpetuity, but it also typically involves an annual support cost. The tool defaults to 18% but is user adjustable in the Advanced Inputs. Choosing "perpetual license" will calculate a 5-year cashflow that shows a payback period for the investment. The subscription cost model, on the other hand, involves paying either a monthly or yearly fee to use the software without any added support cost. When a subscription cost model is chosen in the tool, it will calculate a 5-year cash flow that shows an ROI percentage.

A chart is provided for cumulative cashflows of the DCIM solution cost and value (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 DCIM solution over the 5 years. For a perpetual license-based solution, this equals the one-time license cost (incurred at year 0) and the annual support/maintenance contract cost (default value of 18% per year) starting at year 1. For a subscription-based solution, the cost equals the annualized subscription price (may be incurred monthly or annually). There is no added or separate support/maintenance contract with subscription-based solutions.

The green line represents the value that the use of the DCIM solution is expected to provide based on the inputs and assumptions used in the model. It is a summation of the downtime events avoided, a reduction in staff time spent managing and maintaining infrastructure devices, and the value of eliminating security and environmental events that might occur with the related DCIM monitoring functions in use. The impact of DCIM is summarized by calculating a Return on Investment (ROI) in the case of the subscription cost model and a payback period for the perpetual license cost model.

A 5-year summary of DCIM value is also provided to demonstrate the cumulative downtime costs, staff costs, and security/environmental incident costs (if selected) with and without DCIM monitoring implemented. Having this side-by-side comparison of each category helps illustrate where DCIM has the most impact for the scenario entered.

We ran three scenarios of a distributed IT environment with 100 power and cooling devices, to illustrate how the value of DCIM monitoring can vary, and what some of those key drivers are. **Figure 4** presents these scenarios and shows how cost of downtime, cost of staffing/servicing, and frequency of avoidable incidents can affect the ROI or payback. See **Appendix** for further details of each scenario.

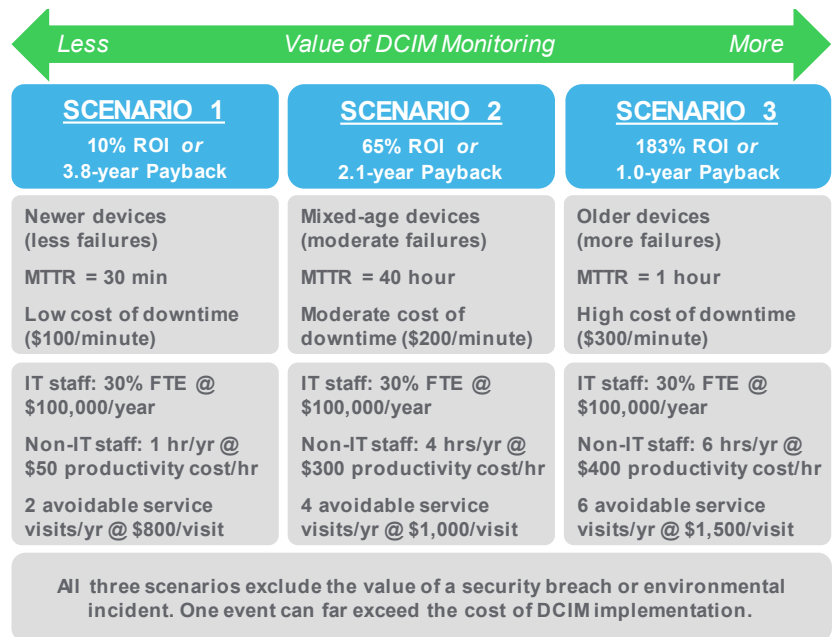


Figure 4
Three scenarios showing how DCIM monitoring value varies based on downtime and staffing baseline conditions

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

Conclusion

Edge sites are critical to business operations, and maintaining availability is essential. But these increasingly distributed, hybrid IT environments often have little to no onsite staffing, making it even more important to have remote visibility to infrastructure device health and environmental conditions. Modern DCIM software provides useful analytics, pre-event failure warnings and useful recommendations. New DCIM is also designed and priced for distributed edge applications vs. legacy software for large standalone data centers. Distributed sites can be very expensive to service due to travel distance. Additionally, outage duration can be extended due to this travel time.

But justifying the investment in DCIM monitoring with an ROI or payback analysis can be difficult to do in a credible way. In this paper, we provide a value framework that includes 1) a reduction in downtime occurrences, 2) a reduction in staff inefficiencies and service costs related to power and cooling device management, and 3) a reduction in security and environmental incidents.

CIOs and their IT operations teams can use the newly developed [TradeOff Tool](#) to quantify the value of deploying a modern DCIM edge tool, based on their individual cost of downtime, cost of service, projected number of incidents, and other user-specific attributes.

About the authors

Patrick Donovan is a Senior Research Analyst for the Energy Management Research Center at Schneider Electric. He has over 27 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





[How Modern DCIM Addresses CIO Management Challenges within Distributed Hybrid IT Environments](#)

White Paper 281



[Cybersecurity Guidance for Data Center Power and Cooling Infrastructure Systems](#)

White Paper 216



[Browse all white papers](#)

whitepapers.apc.com



[DCIM Monitoring Value Calculator for Distributed IT](#)

TradeOff Tool 29



[Browse all TradeOff Tools™](#)

tools.apc.com



[Learn more about the next evolution of DCIM](#)

se.com/dcim

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

Appendix

Three screenshots of the [TradeOff Tool](#) are shown below to provide detail on the scenarios presenting varying ROI and payback, as summarized earlier in **Figure 4**.

Figure A1

Scenario 1, showing ROI of 10% with subscription-based DCIM (or less than 4-year payback with perpetual license-based DCIM)

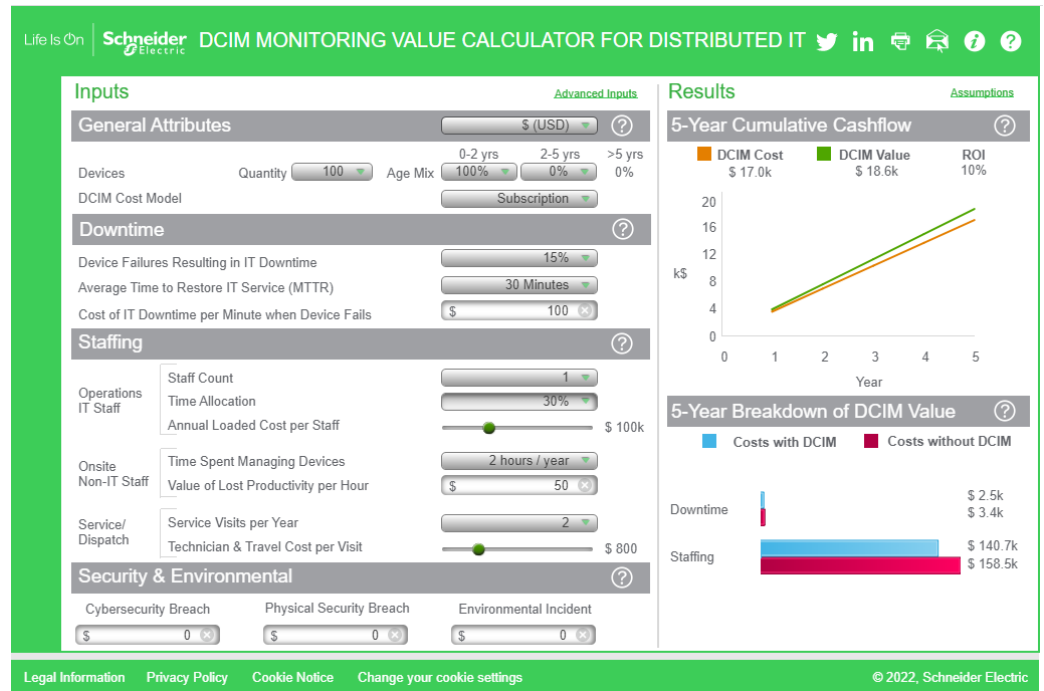


Figure A2

Scenario 2, showing ROI of 65% with subscription-based DCIM (or approximately a 2-year payback with perpetual license-based DCIM)




Figure A3

Scenario 3, showing ROI of 183% with subscription-based DCIM (or a 1-year payback with perpetual license-based DCIM)

