Understanding Multicloud Monitoring: Ensure Success the First Time

David S. Linthicum

Executive Summary

Let's face it, the world of IT is getting more complex. Many people view private and public clouds as something that will replace traditional computing and eventually make things easier. However, for most enterprises, the cloud is now an additive to existing traditional systems, and that adds complexity.

There's also the need to build net-new cloud-native applications for the business. The result is that more silos are being stood up while the complexity and the heterogeneity of the platforms continues to increase. The accepted fact that all enterprises will end up with multicloud deployments is now the reality within organizations that are already on the cloud journey. Consider a recent study "HYBRID CLOUD USAGE POSES NEW CHALLENGES FOR MONITORING SOLUTIONS" by Dimensional Research. In a Global Survey of Cloud Professionals, the following cloud adoption trends where noted: 92% use multiple cloud vendors already, 88% of cloud-based applications share data and services with on-premise applications, 64% indicate cloud-based apps share data and services with other cloud-based apps.

But does multicloud mean even more complexity?

No. The multicloud environment will be debilitating only if we choose to make it so. If employed correctly, emerging approaches and tools will have the ability to solve today's complexity problems, and eliminate them moving forward.

That is, if you're willing to learn how. Reading this paper is a step in the right direction.

Forrester stated in 2016 that, while organizations already use multiple clouds, they will do so even more so in 2017. Why? CIOs continue to step up to orchestrate the various cloud ecosystems that connect employees, customers, partners, vendors and – with the Internet of Things (IoT) in mind – devices to serve rising customer expectations.

Also, according to Forrester's predictions, enterprises won't just have an increasingly multi-cloud reality to orchestrate. They will become cloud companies themselves to provide core cloud services to their customers and partners.

Another analysts group that discovered the shift to multicloud is OVUM. Their report, "Realising possibilities in the cloud: The need for a trusted broker," focuses on a fundamental shift in cloud purchasing and mentions how more than 40% of European organizations and enterprises expect to manage a hybrid multicloud environment within the next two years.

The number of data points are starting to grow.

Proactive monitoring is key

This paper presents several approaches to multicloud monitoring, including native and non-native tools.

Native cloud monitoring tools leverage systems furnished by the public or private cloud provider. You leverage tools such as AWS's CloudWatch, or other tools that are a part of the cloud services portfolio of the specific public cloud provider.

The tradeoff here is pretty clear. If you leverage a single cloud provider, this would be the lowest cost option and the easiest to deploy. However, we're clearly moving to a multicloud world, where we leverage more than a single public cloud provider. Monitoring more than one public cloud using native monitoring tools means we'll have to use more than a single monitoring tool as well

This paper comes to the conclusion that this multi-tool approach is obviously not optimal. We have to learn more than one tool, monitoring quickly gets confusing, and the opportunities increase to make mistakes or miss data points. This has largely been proven to be an ineffective approach, or, at the very least, much less effective than more comprehensive tools that can provide a "single pane of glass" with visibility into all public and private clouds under management.

Non-native tools are perhaps the best approach. You leverage a single tool to monitor several public and private clouds. They also have the capability to monitor traditional platforms as well, including legacy, ERP, distributed computing, network storage, and other resources that might be a part of your on-premises solution.

Conclusions and recommendations

After reading this paper, you should understand why we make the following recommendations, as well as understand the basic approaches for following these recommendations:

Use monitoring tools that are designed for cloud heterogeneity. A single tool that's purpose-built for a single cloud provider won't do much when your enterprise is likely to leverage several public and private clouds, now or in the near future. Heterogeneity is on the critical path. At the end of the public cloud journey, we're likely to end up with a mix of legacy systems that remain on premises, as well as new IoT devices and other mechanisms that exist in the domain of IT. We'll also have a mix of private and public clouds, plural, and the need to have all of the above share information and core business processes. We need technology to help us get there.

Leverage a configuration management database (CMDB). The CMDB is a central point of your efforts that allows you to see your IT assets at many levels of granularity. A CMDB represents the authorized configuration of the significant components of the IT environment, and is a must-have to do multicloud and traditional system monitoring.

Have a mechanism to sense, analyze, adapt, and visualize which will help admins solve key tactical problems before they become outages. To start, you need to sense (find) everything in your environment. AI Operations will analyze all the data coming from what you've found and activate a runbook if needed. A typical runbook contains procedures to begin, stop, supervise, and debug the system, and may also describe procedures to handle special requests and contingencies. Apply changes to the environment based upon what you sense and analyze within the runbooks. You need to provide visualization to gain a holistic view of the runbook. All features are required to support a multicloud deployment.

Leverage a monitoring tool that provides anomaly detections across all cloud and traditional platforms, that can also treat each event the same. This centers around the actions required to identify events or observations which do not conform to an expected pattern or other items in a dataset, or in the history of data gathered while monitoring the systems.

Leverage monitoring tools that provide automation. The ability to automate the functions of the tools allows you to do more with less, as well as create automated procedures that keep your systems running without human intervention.

Leverage monitoring analytics. Gather data over time for the targets, and then make tactical and strategic sense from the data. This can be bound with automation to determine trends that could be leading to outages, and have the monitoring tool take automated corrective action.

To understand why we make these recommendations, please read on.

Introduction

The adoption of cloud computing is quickly moving to a multicloud practice. Not long ago, we only envisioned a single public cloud provider and a few workloads. Now we see several providers, each running localized workloads. To complicate matters, most of these instances communicate with other application workloads, either intra-cloud (same cloud provider), inter-cloud (different cloud providers), or with traditional systems or private cloud on premises.

This intermingling of systems, which are not coupled to a platform (e.g., a public cloud provider), is just the logical conclusion to the use of these distributed systems. Leveraging the right platforms for the right applications provides the ultimate value and justification for leveraging cloud in the first place.

However, this also leads to the challenges of complexity, and the ability to make complex systems work and play well together in the long-term. In other words, when we plan for cloud computing in the enterprise, we need to understand that complexity will be a prominent feature of the end state. The first step to prepare for that complexity is to get the right ops processes and tools in place, which include management and monitoring tools. During the remainder of the paper, we'll tell you how to complete this step, and what issues and concepts to consider.

The challenges of cloud monitoring

So, where are we likely to progress? Figure 1 shows the likely growth of both cloud and on-premises systems over the next 29 months, and the linked rise of complexity. Considering this data, here's what we can learn:

- The inflection point won't provide us with much reaction time. Note in the chart that the enterprises will have roughly 6-8 months to get their monitoring and management act together, starting from the deployment onto public clouds. For most enterprises, that's not enough time. Moreover, you're dealing with technology that already supports production.
- Considering that complexity will inflect as well at the time the number of systems do, cloud or not, we must deal with three growth curves, not two or one. This will drive additional needs such as security, governance, cost management, etc., which need to be done at the same time as monitoring.

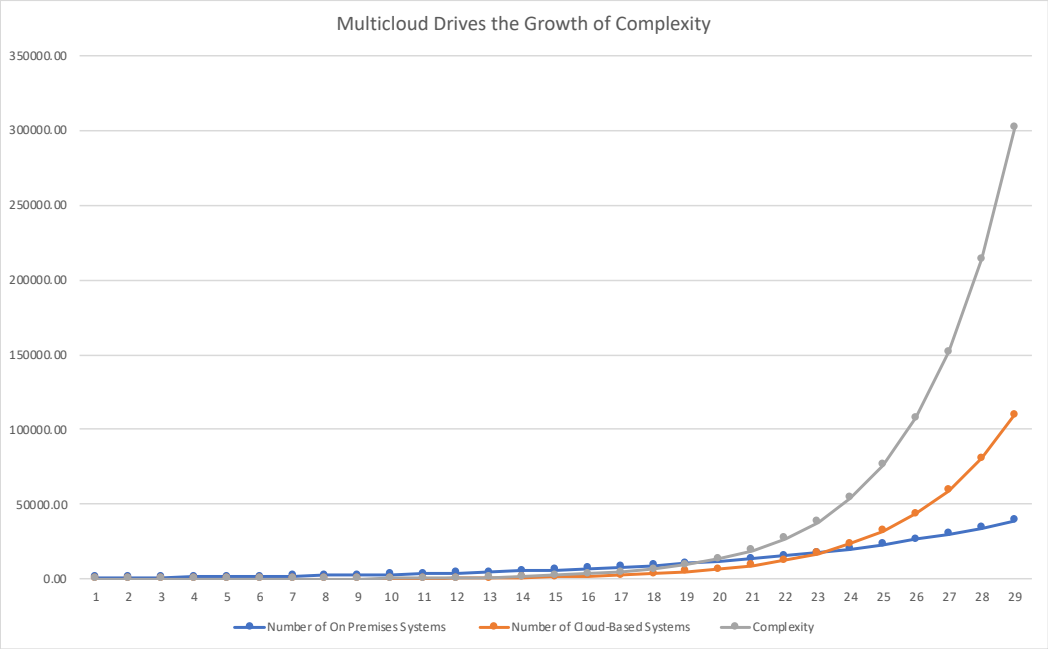

Figure 1: This chart covers 29 months, as a sample of time from the beginning of cloud planning, throughout deployment, and then with focus on operations. It shows the growth of on-premises systems and cloud-based systems, as well as the effect on complexity. (Source: Linthicum Research, LLC.)

The path of least resistance is to leverage the native monitoring tools found within clouds, for example, AWS CloudWatch. Considering the chart in Figure 1, it may be possible to monitor the systems months 1 – 15 using the cloud-native monitoring tools. However, those will likely hit their tipping point in terms of complexity between months 10-17. This means that the number of on-premises systems, as well as cloud-based systems, begins to grow faster than the operators can easily manage using native tools. Because they need to use one tool for one cloud, within multicloud situations their ability to monitoring clouds using 2 or more separate tools will eventually hit a tipping point. What's bad about this tipping point is that the lead-up symptoms will include system failures caused by the inability to properly monitor the systems.

Moreover, and as important, is the need to monitor the complete end-to-end path of an application to quickly determine the point of failure. As complexity goes up, we also have to deal with system dependencies that cause issues. We need the ability to find and resolve those problems before they become issues, or, after they become issues and are spotted by the operators.

It comes down to determining the point of failure within complexity. We need the ability to see the application failure, as well as the root cause of the failure. That means being able to determine anomalies and adverse trends to proactively identify issues before actual incidents disrupt business

For instance, if a cloud-based application begins to show performance problems, we need to determine the root cause of the problem. While many would look to saturation of the virtual

compute server in the cloud where the application is hosted, other issues are actually more likely. In this case, the operator determined that a network switch was misconfigured in the on-premises network. This was determined by monitoring the application, the cloud resource where it's running, and the underlying network which is taking the application interface from the application on the cloud provider down to the user through the network.

In this scenario, the cloud-native tool could only monitor the cloud resource that hosts the application, so you would have had no clue as to the root cause of the issues. Multiclouds, including multiclouds mixed with on-premises systems, are complex distributed systems. Applications are built with many interdependencies. You need to monitor the systems on their own, as well as for dependency. Keep this in mind: Systems that span many different clouds, and perhaps systems based on-premises, need to be monitored with an understanding of these dependencies. This is where it's critical to have a CMDB. Otherwise they may work on their own, but fail as a collective whole. This makes monitoring and management that much more complex, and emphasizes the need for automated monitoring tools to assist in solving this problem.

Furthermore, you need a way to express the state of applications in terms the business can understand. For example, the database should be named customers, transactions, inventory, etc., and not DB1, DB2, DB3. DB3 going offline for no reason could confuse an operator who's managing hundreds of databases and something else might get 'fixed.'

Additionally, you need a way to express the state of applications in terms the business can understand. This could require many operators at different levels of experience to understand just what's going on, and have them put that understanding in writing via runbooks. Common nomenclature and models that represent these objects are required or people will refer to the same resource with different legends. It isn't enough to uniquely identify and name an object, understand the dependencies that exist between them is also required.

Cost monitoring is a consideration as well, along with system metrics. While a good monitoring process is half of the battle, you also need ties into cost metrics and governance.

All activities that occur on private or public clouds need to be tracked. What are the processing needs of deployed applications in the multicloud, and how well are they keeping up? What's being spent in terms of cloud costs, both internal costs, and those that will be billed by the public cloud provider? You need those answers at your fingertips.

There are a few layers to the answers. First, you need the ability to monitor most data points on all clouds under management. Second, you need the ability to gather data and metrics about the ongoing costs of activities. This becomes complex unto itself, considering how public cloud providers bill for usage. Sometimes it's by the minute, sometimes by the hour, sometimes by the amount of usage (such as GBs stored), and sometimes as a fixed fee for a block of resources.

This problem extends to the business issues as well. We need to monitor the activities of the multicloud to make sure each cloud does their part to insure the health of the overall applications, we also have to track the ongoing costs of all this activity.

It's not only about monitoring costs and usage by resources, users, and other types of metrics you may need to use. It's also about your ability to place limitations on what can be spent. The days of getting $1M cloud bills are long over, and businesses look at public clouds as any other type of utility; something that needs to be monitored, budgeted, and governed ongoing. Doing this directly from your monitoring solutions seems like a logical step.

Cost impact of complexity

So, what if this disease goes untreated? As you can see in Figure 2 below, the costs inflect as well, to a degree that most enterprises cannot afford. This assumes no plan for multicloud monitoring, no tools employed, and no regard for or reaction to increased growth and complexity.

Enterprises see diminished value from cloud computing due to these issues. In fact, cloud computing may not make sense for organizations that do not address the complexity that cloud and multicloud will bring to enterprise IT.

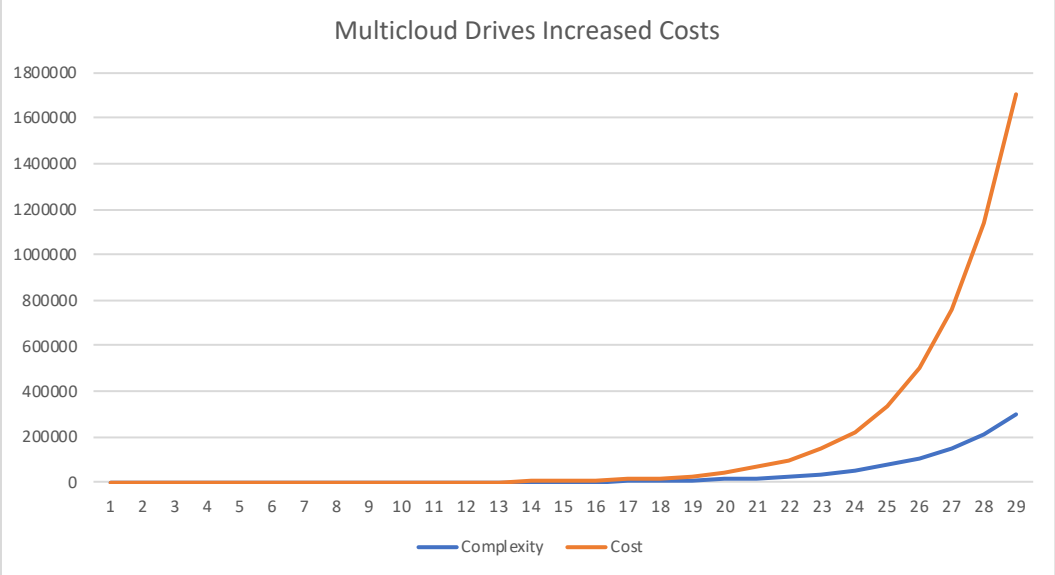

Figure 2: As multicloud complexity inflects, so does the cost - well beyond budget. (Source: Linthicum Research, LLC.)

Moreover, there is a people impact as complexity increases, and that typically translates into more money and risk. For instance, there are the aspects of different cloud stacks to consider and the complexity when you move into multicloud, such as hiring experts for the specific cloud brands, and dealing with different cloud-native features, such as databases, machine learning, and native monitoring systems.

A better approach would be to automate the management and monitoring of those systems. This creates an abstraction layer between the complexity and the people. Along with generalized automation (discussed next), automated systems management and monitoring should result in the need for fewer humans to manage the multicloud, and a much better operational experience.

Another thing to consider are the requirements when cloud-native apps are being developed, and the need to customize the configurations within the applications (also called "monitoring as code"). You need to provide monitoring configuration information within the application itself, or have the application tell the monitoring system how it wants to be monitored on the cloud, as well as the configuration that's best for the application.

Strategic vs. tactical thinking

We've made the case for the limitations of cloud-native monitoring. Now it's a good idea to break down the core limitations. They include:

Operations – No single pane of glass to view all resources, on all clouds under management. The use of multiple consoles is never a good idea. It creates confusion, lacks visibility into all systems, lacks the ability to create rules and policies across all cloud providers.

Outage – No core insight into what went wrong, how to prevent it in the future, nor how to spot system behavior that leads up to an outage. We all understand that monitoring is needed to prevent things from going wrong, but if you lack the necessary visibility and data points, you'll be courting outages. Internal and external.

Security – No holistic visibility means that we can't spot attacks, and take evasive action. The core tenant of security is to have visibility across all systems that can be attacked, and the ability to watch for unusual behaviors or trends that lead up to an attack pattern.

Customer impact – No information about the impact of outages, security issues, and any operational issues on the customer experience. While some negative system event may only cost a few hours of lost productivity, the negative impact on the customer could cost you many times what you lost in the internal event. Customers these days expect zero downtime. More and more, they vote with their dollars.

Automation and Analytics

Automation means automating the configuration of the monitoring for all infrastructure and application components. While the approaches differ, the use of automation deploys the right monitoring configuration to the target systems that are under management. (Many use agent-based or agentless monitoring approaches.) You also need the ability to detect new instances and automatically include them into the system monitoring groups. Use prebuilt templates to properly configure them.

Automation means that we have a more holistic management solution because we can deal with issues using automation, and not have to toss people at the problem. Once we figure out a path to solve a repeatable problem, automation can solve that problem for good. Also keep in mind that automation allows us to better respond to change, such as the ability to dynamically deal with changes, and allow the monitoring solution to operate independent of the specific cloud brand or traditional system under management.

In the context of machine learning, 'analytics' is the ability to gather information ongoing from target systems, storing that information for access, and then running all types of analytics using the data to determine different things around the systems that are being monitored.

This seems innocuous, if you just look at status or performance data coming from a database server. However, the idea is to look at patterns within the data, such as performance issues that lead up to network failures, or databases that produce increasing I/O errors. These patterns allow us to spot and correct issues before they become issues.

In the world of management and monitoring there are tactical aspects to analytics. The real value is more strategic. In other words, look at the health of specific systems, but understand the bigger pictures as well, such as performance issues that could be leading up to more public cloud costs by leveraging more instances, or capitalize on the ability to see dependencies between systems within complex migrations that are causing issues that need to be corrected.

Core tenants of multicloud monitoring

We defined the problem in this paper, and the impact of the problem. Now, let's define the solution. As seen in Figure 3, there are four core tenants required for an optimal multicloud monitoring solution:

Leverage a higher level of abstraction. Abstract complexity from the operational user. Provide the humans who operate the multicloud system with the ability to have views into the systems that are most relevant to and helpful for their needs. What they see and how they automate the processes should be removed from the complexities of the underlying system. They can't track down to all levels of detail, thus the abstraction layer shows them just what is relevant.

Monitor heterogeneity using a single pane of glass. Provide a console that spans many systems. Use the layer of abstraction, automation, and data analytics to provide just the relevant information that pertains to the health of all cloud systems under management.

Monitor analytics across clouds. Gather data from all clouds under management, ongoing. Analyze the data by cloud and between clouds, and by applications and between applications, to find issues and trends before they impact the business.

Proactively monitor across clouds. Find issues before they happen. Have the ability to monitor trends, such as poorly preforming storage instances or other issues that could lead to networking and system outages. The tool's ability to be proactive leads to the maximum uptime and protects the business, which should be the clearly stated objective of the operators who use the monitoring tools.

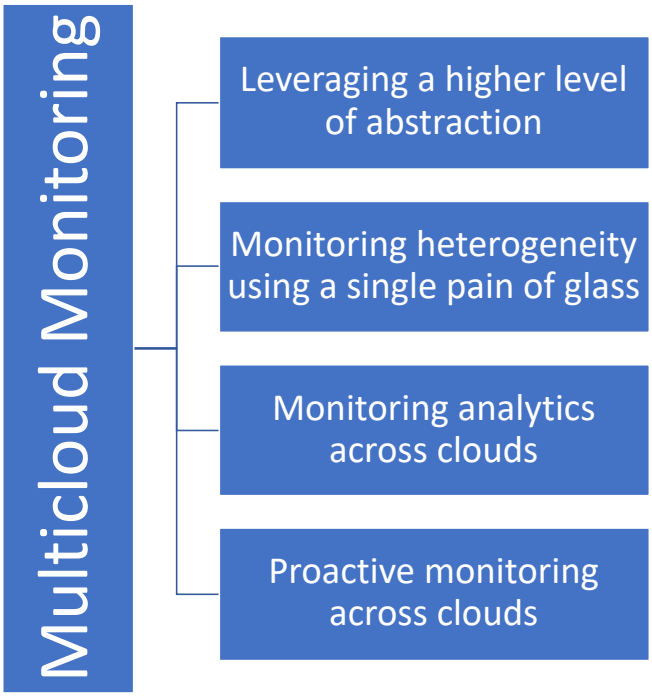

Figure 3: The four core tenants of multicloud monitoring.

Leveraging the right technology

So, now that we understand the problems and understand the solution patterns, how do we pick the right tools for the job? It's a matter of defining your criteria, looking at the market, and picking the right solution or solutions.

Let's define some general technology criteria for picking multicloud monitoring tools:

Criteria	What to look for	Scoring
Cloud heterogeneity	The ability to monitor the greatest number of private and public clouds. Tools must dynamically adapt to changes at minimal cost. They also need to optimize allocations and negotiate/renegotiate cloud services as usage increases.	Does the monitoring tool connect to the major cloud brands, and, if so, does it leverage all points of monitoring?
Monitoring and analytics	The ability to generate reports, and dashboards that are most relevant to those who operate the multicloud. Monitoring solution must exploit the analysis of multiple types of data and understand how to leverage that data in aggregate using different layers of abstraction.	Does it provide the reports and dashboards required by the operators? Does it provide custom reports, and custom dashboards? Can the analytics be proactive enough to spot issues as trending? Can analytics work and play well with other parts of the monitoring tool?
Automation	The ability to automate monitoring and management. This includes proactive capabilities that can solve problems by deploying sequenced and preprogrammed steps.	Does the tool provide enough automation capabilities to satisfy the needs of the stakeholders?
Actionable information	Provide crucial insight to stakeholders that enables them to make better decisions faster and helps IT itself present the value of their work and data in meaningful forms to the stakeholders.	Does this tool provide the information needed by the stakeholders?
Being proactive	The ability to spot trends and provide other data to detect anomalies and predict issues that can be proactively corrected.	The degree that the tools can gather and monitor data that's most relevant to the health of the multicloud, as well as generate alerts before the issues become critical.
Self-healing	The ability to spot issues and take automated corrective action.	The amount of self-healing capabilities relevant to your multicloud solution.
Evolving over time (agility)	The ability for the tool to change over time as the needs of the business change, as well as the technology solution.	How well can you expand or contract the monitoring solution to add or remove technology needed by the business?
Cost effectiveness	The tool's ability to save money beyond its cost.	ROI
CMDB	Make sure to leverage a configuration management database (CMDB).	CMDB meeting business expectations?

Table 1: Criteria for selecting and scoring multicloud management tools.

Call to action

The call to action is to get your multicloud monitoring strategy underway now, or suffer the costs and business impact consequences. The good news is that the value of multicloud far exceeds its cost and risk, that is, if you can get your arms around the multicloud monitoring requirements right now rather than wait for the complexity inflection points outlined above.

Automation is on the critical path. Keep in mind that the increased complexity and the increased data volumes generated will stress the IT Department, often at the same time their budgets are reduced. Look at automation as a way to avoid hiring more humans, and as a way to provide better systems monitoring, management, and thus uptime.

There is also the issue of compliance. Compliance can be linked to service delivery with SLA compliancy, or linked to corporate compliance that shows how corporate policies are respected, or national/international compliance according to rules and regulations. Automation provides a clear methodology to present and prove consistent behavior of the IT Operations activities at minimal cost.

It doesn't matter if you're moving to a complex multicloud that consists of many public and private cloud computing brands, or merely a single hybrid cloud. Monitoring is core to what makes this technology work for the business. While some enterprises will put off the use of heterogeneous multicloud monitoring for as long as possible, they'll find that a quick catchup is impossible. When the complexity inflects, you'll end up failing, and perhaps take down the business as well.

The best analogy that one can think of is that of monitoring a power plant. While we can certainly leverage the gauges that came with all of the equipment to monitor the basic workings of the plant, generators, boilers, etc., there is no way you can figure out the abstract state of those hundreds of gauges and how they are behaving in terms of a normal day, or on a day when the whole plant could blow up.

Multicloud monitoring tools put critical information in the right hands at the right time so decisions can be made by both humans and automation. Moreover, this is about the collection of information and understanding how that information relates to itself and each other, and it's true meaning as it relates to the overall health of the multicloud system. Most importantly, these tools tell us how we can keep things running at the lowest cost, with the lowest risk.

There are no magic pills that will eliminate your multicloud issues, no matter how good your monitoring solution. The goal is to reduce the number of issues that come about on a daily basis, and insure the uninterrupted use of multicloud to solve problems for the business on a consistent and reliable basis. It's all about making the business successful. That's the core metric you need to pay attention to.

About the author

David Linthicum was just named the #1 cloud influencer via a recent major report by Apollo Research. David is a cloud computing thought leader, executive, consultant, author, and speaker. He has been a CTO five times for both public and private companies, and a CEO two times in the last 25 years.

With more than 13 books on computing, more than 5,000 published articles, more than 500 conference presentations and numerous appearances on radio and TV programs, David has spent the last 20 years leading, showing, and teaching businesses how to use resources more productively and constantly innovate. He has expanded the vision of both startups and established corporations as to what is possible and achievable.

David is a Gigaom research analyst and writes prolifically for TechBeacon, and for InfoWorld as a cloud computing blogger. David is also a contributor to "IEEE Cloud Computing," Tech Target's SearchCloud and SearchAWS, and is quoted in major business publications that include Forbes, Business Week, The Wall Street Journal, and the LA Times. David has appeared on NPR several times as a computing industry commentator, and does a weekly podcast on cloud computing.