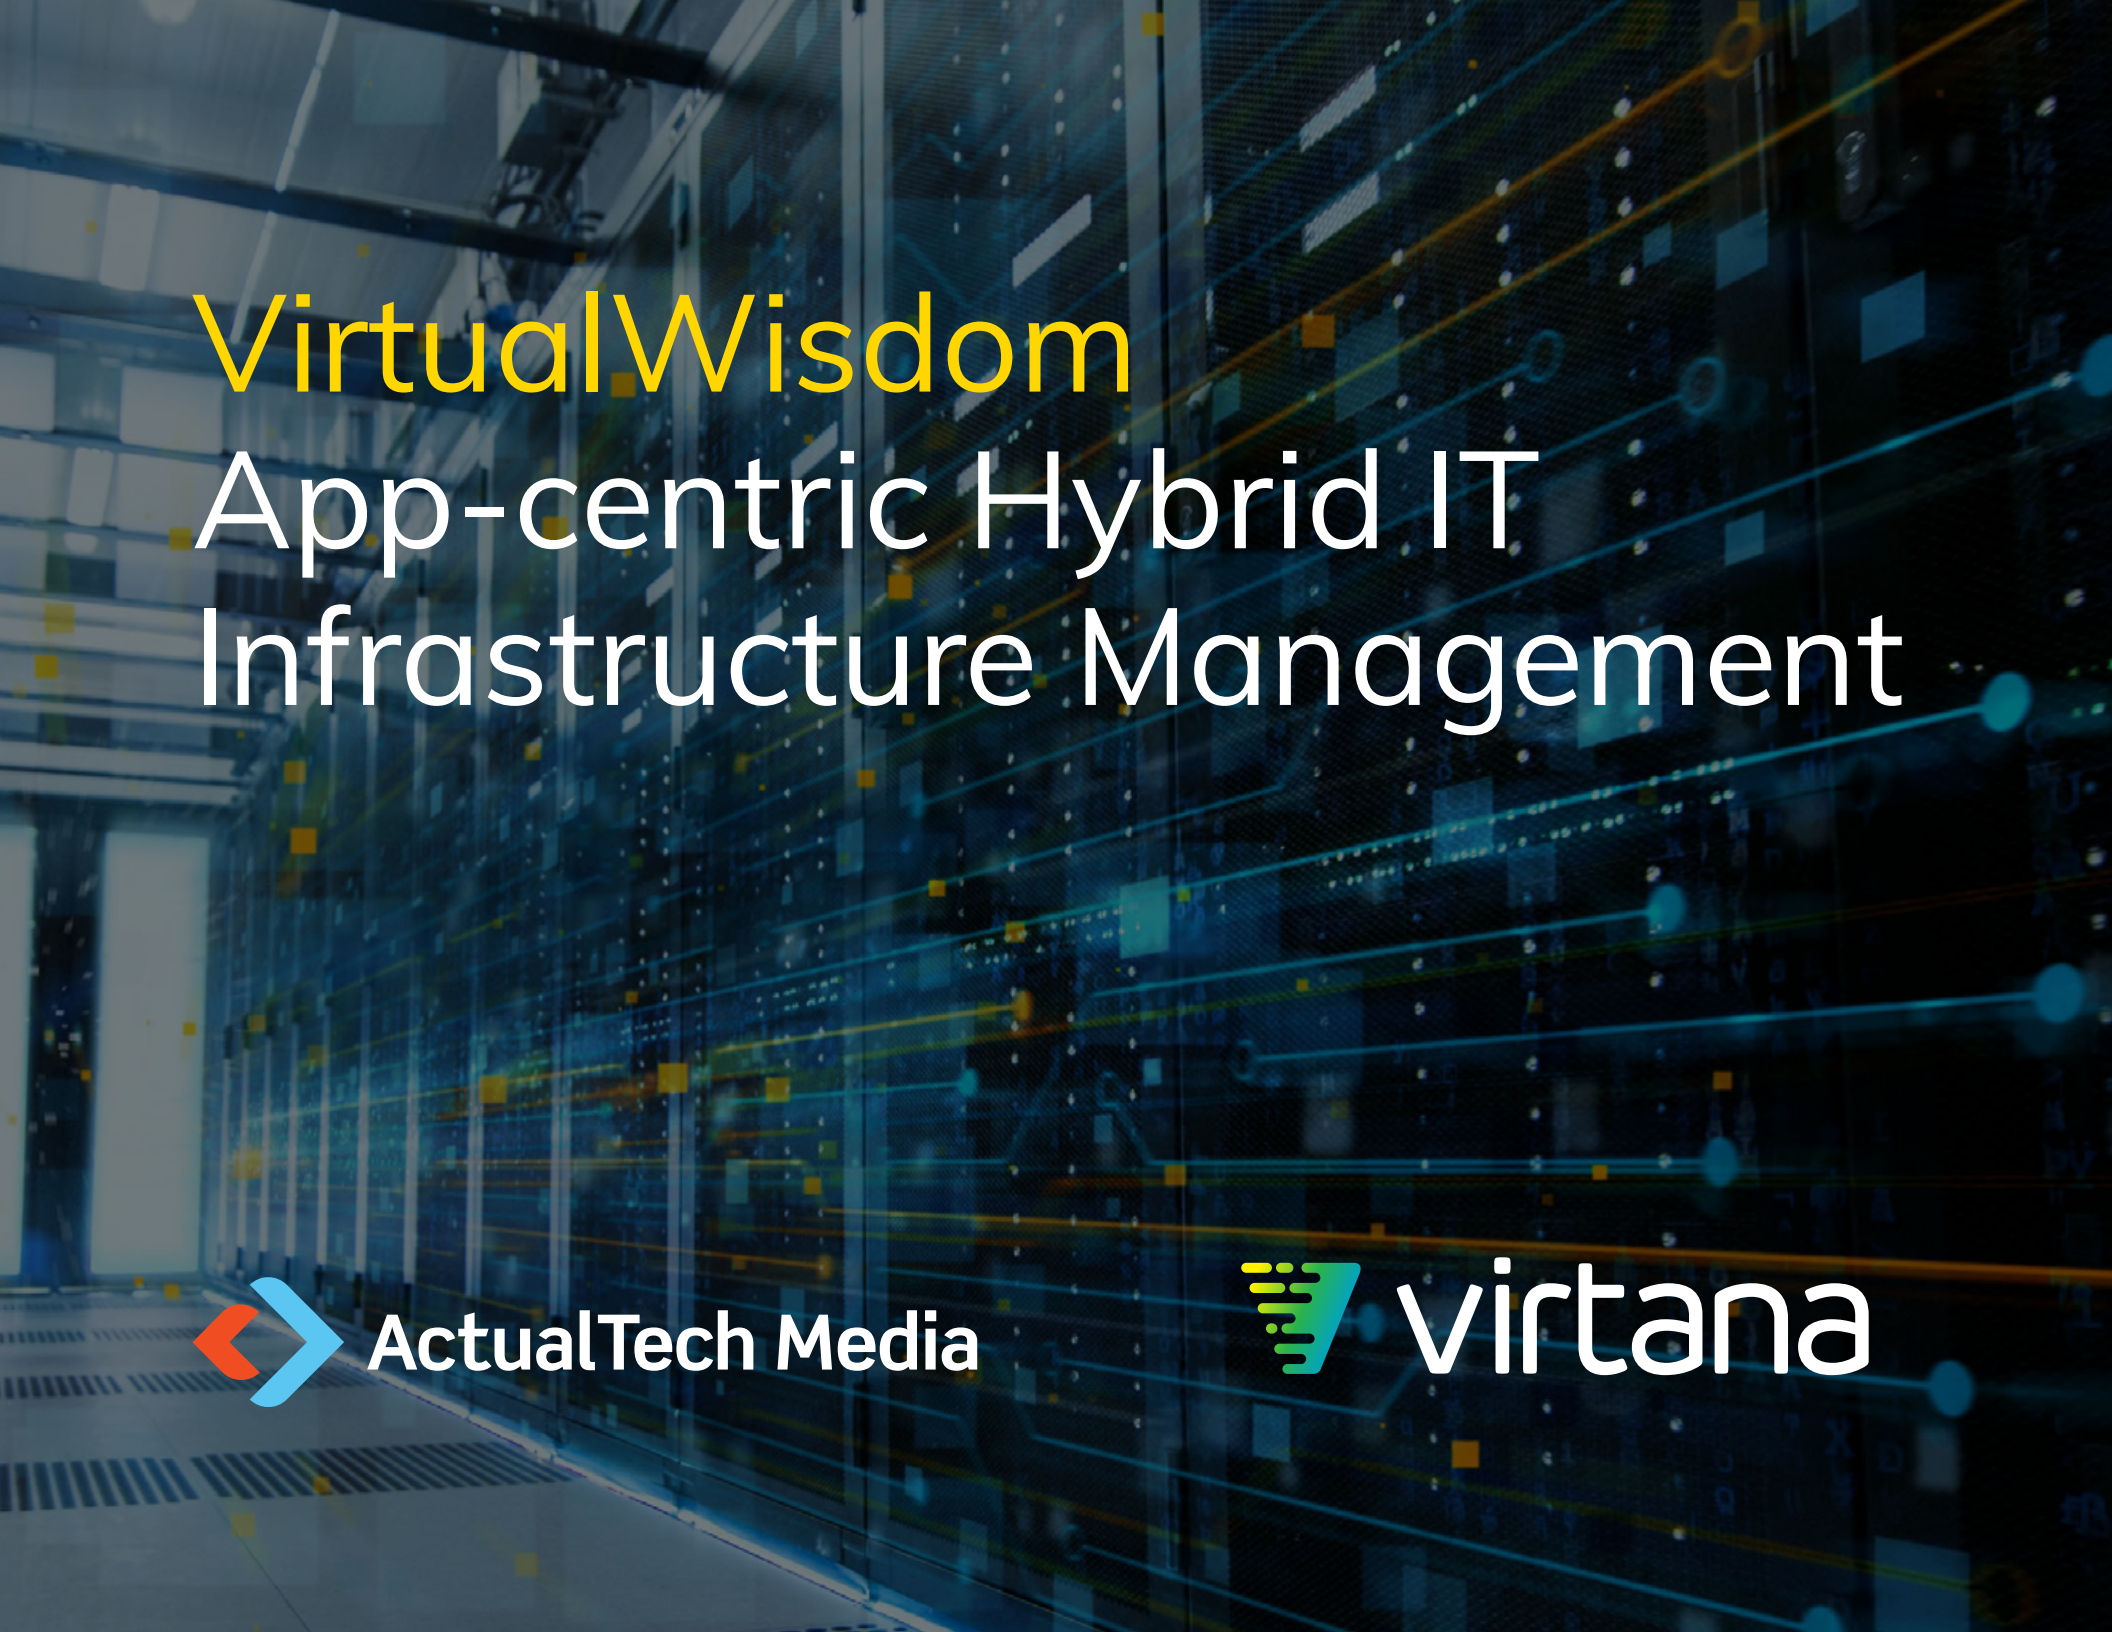
VirtualWisdom
App-centric Hybrid IT
Infrastructure Management
ActualTech Media
virtana

Contents

Introduction

Slow application performance and application outages cost money. As organizations have become more dependent on IT, the cost of outages has risen. Reducing the number of outages requires identifying their cause before the business is impacted. This is exactly what VirtualWisdom App-centric Hybrid IT Infrastructure Management excels at.

Today, the cost of downtime for a single server can be punitive. According to **Statista**, 24 percent of respondents worldwide reported the average hourly downtime cost of their servers as being between $300,000 - $400,000, with more than 50 percent of respondents reporting costs in higher brackets.

These statistics echo those released by **Gartner**, who cited $5,600 per minute as the cost of downtime: $5,600 per minute is $336,000 per hour.

In today's highly consolidated data centers, multiple individual virtualized hosts typically share a common infrastructure. A failure in any component of that infrastructure could result in multiple virtualized host failures. While it's difficult to determine how much more a multiple-server outage would cost a given organization over a single server outage, the answer is certainly more than a single server outage would cost, and those costs are already far too high.

OPERATIONAL VULNERABILITY

The cost of outages is magnified for organizations that have service level agreements (SLAs) with their customers for IT availability. Traditionally, SLAs have been the province of hosting service providers and cloud service providers; however, this is changing.

Organizations of all sizes increasingly make critical information available to their customers and suppliers through both on-premises and cloud-hosted systems. An example would be a Financial Accounts processing business that makes its accounts system visible to customers. Customers quickly become familiar with the data made available to them through the accounts system, in many cases building that data into their own internal IT and business processes, sometimes incorporating that data in the information they themselves reveal to their customers.

The Financial Accounts processing business might also be sharing data with other brokers. market trading systems are often automated, even for smaller institutions. Similarly, integration of order tracking, billing systems, and logistics systems is crucial for most organizations that transact in volume.

Here, our example of a Financial Accounts business might have both availability and performance SLAs with both customers and market institutions. Not meeting SLAs with means that necessary data for trading may not arrive on time, and/or transactions may not complete on time. Not meeting SLAs with customers may also require the business to pay penalties to customers, too.

Not meeting market institution SLAs means that necessary data for trading may not arrive on time, and/or transactions may not complete on time. Not meeting SLAs may also require the business to pay penalties to customers, too.

IT'S NOT GETTING EASIER

Applications rarely exist on an infrastructural island in today's data centers. Most applications exist in a virtual machine or container, which executes atop a workload supervisor (i.e., a hypervisor or microvisor), which in turn runs atop a physical compute server but likely uses shared storage and participates in a compute cluster.

Most applications are also reliant on many network services (e.g., firewalls, reverse proxies, and load balancers). A problem with any of these physical or virtual infrastructure components can lead to performance degradation of the application, or even a complete outage.

Resolving application outages requires learning which infrastructure component(s) is/are responsible for the outage. Uncovering which components are responsible for an application's misbehavior typically requires the correlation of information about the application itself and multiple infrastructure components.

Determining which shared infrastructure components an application depends on isn't easy. Correlating information for all those systems is even harder. Even if IT stakeholders know what must be done to identify the root cause of an outage, performing that analysis manually is like finding a needle in the haystack.

Complicating matters is the fact that infrastructure complexity grows with time and the fact that most enterprise-class organizations use multiple (e.g., hybrid/multi-cloud) infrastructures. Multiple infrastructures can make infrastructure data correlation orders of magnitudes more complex, dramatically increasing the size of the haystack through which IT professionals have to sift to find the needle at the root of the outage.

Of course, if you want to find a needle in a haystack efficiently, you don't sift through the hay with your hands. You bring a really big magnet and solve the problem in seconds. VirtualWisdom App-centric Hybrid Infrastructure Management is the magnet: it plucks the root cause of an application outage from the haystack of infrastructure telemetry in seconds.

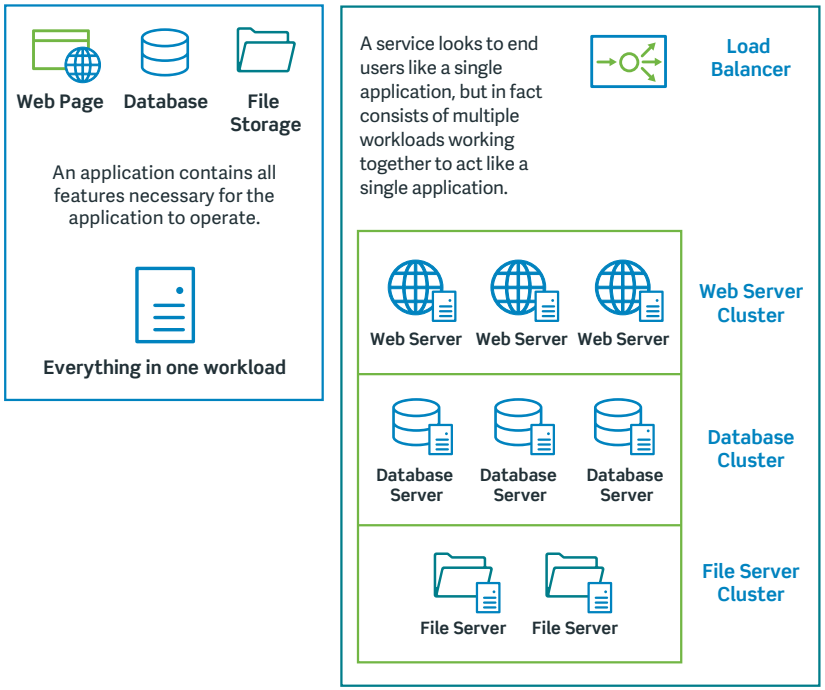

Figure 1: Applications are complicated.

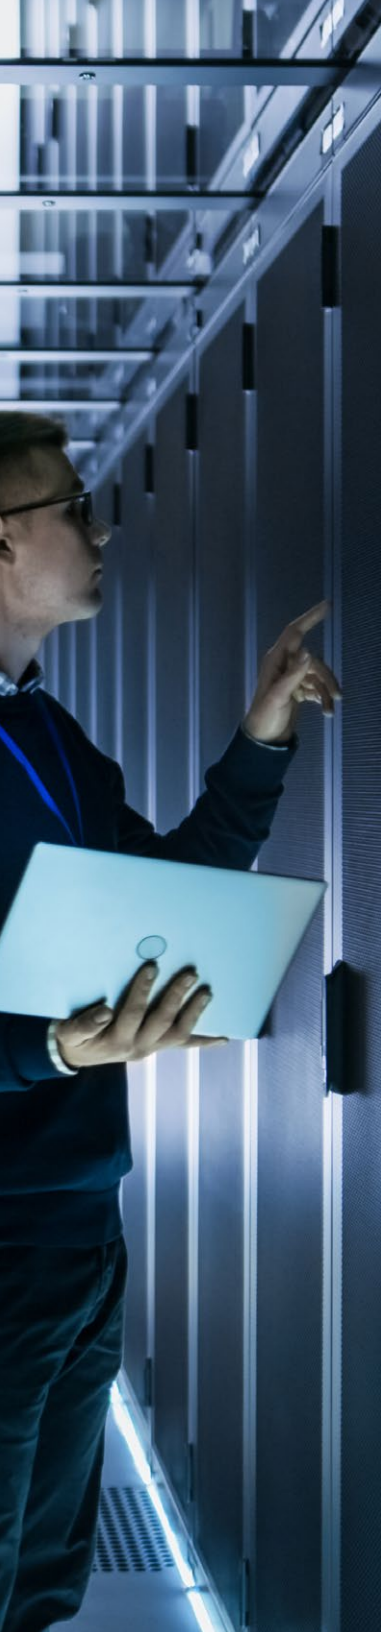

Expertise Required

In addition to outright application outages, application performance degradation is a serious concern. A Vanson Bourne survey of 100 IT managers found that poor application performance has cost a third of the respondents' organizations more than **£1 million** (about $1.2 million U.S.), with 15 percent of respondents reporting that poor application performance has cost them more than £2 million (about $2.4 million U.S.) each year.

The cost of poor application performance is probably much higher than reported, too. Research into the cost of application performance problems is limited because most organizations have wholly inadequate performance tracking for their applications. The Vanson Bourne survey found that 89 percent of organizations had inadequate performance tracking, a figure that's similar in other datasets.

Despite the limited data available on this topic, poor application performance has clear impacts. The **2017 app attention index** (registration required), published by AppDynamics, shows that 80 percent of users delete poorly performing apps. Application users internal to an organization don't have that option, but mounting frustration within a workforce due to inadequate tools is a well-known problem for organizations of all sizes.

A Vanson Bourne survey of 100 IT managers found that poor application performance has cost a third of the respondents' organizations more than £1 million.

WHAT EXPERTISE IS REQUIRED?

Let's consider a hypothetical scenario in which an application admin is responsible for a single corporate application. The purpose of their employment is to monitor and administer this one application; they must respond to any performance issues, solve any outages. They have no other concerns.

While end users see a single "application" that they use and consume, our administrator knows better. What is perceived by end users to be a single application in fact comprises multiple workloads. Starting from the user's side and working backward, there is a user application that connects to a backend infrastructure.

The IT infrastructure—supporting all applications—starts with a firewall cluster, followed by a load balancer cluster, a web server cluster, a file server cluster, and then a database cluster. Each cluster comprises at least two workloads, so that the failure of any one workload can't bring down the whole application.

Each individual workload operates inside a VM, and each VM has its own operating system. Each VM operates on a hypervisor cluster that uses shared storage. Furthermore, all nodes of a hypervisor cluster use multiple redundant network interfaces to connect to one another, to their storage, and to other hypervisor clusters.

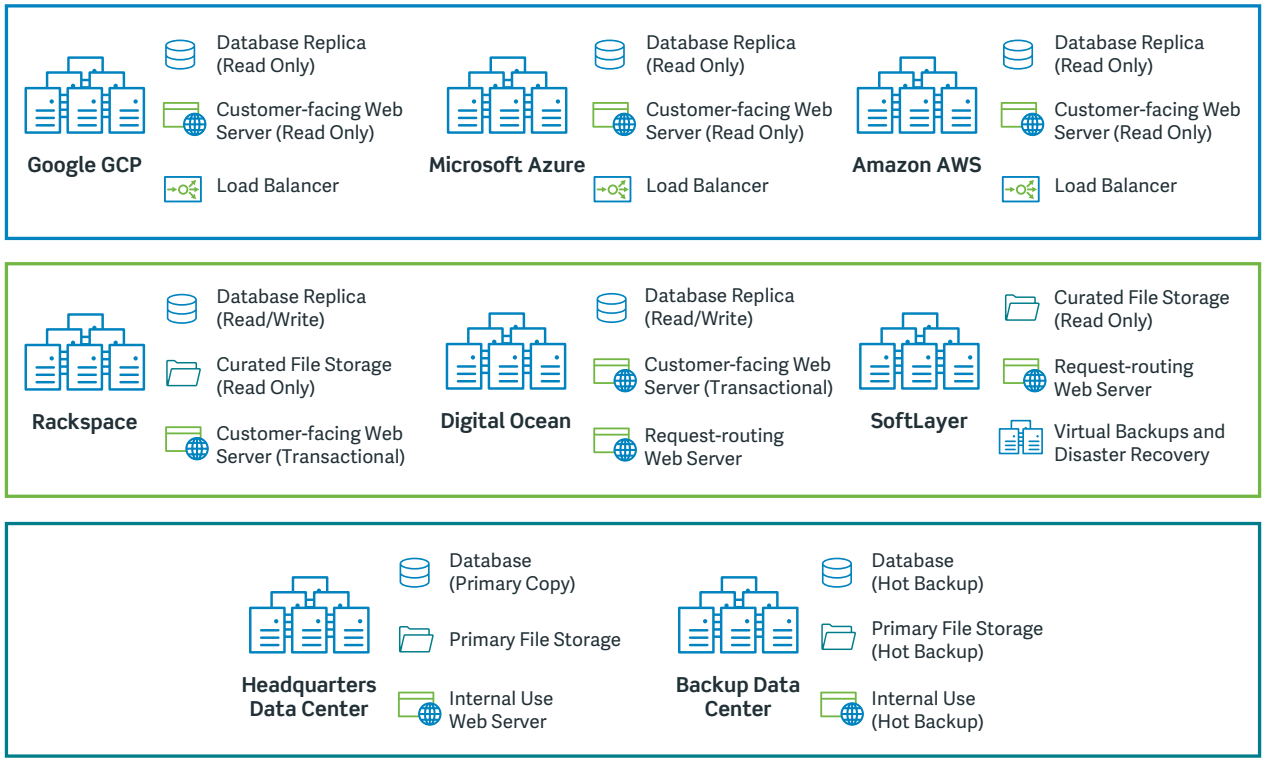

Figure 2: This drawing shows a multi-workload service that presents to end users as a single application, but which is built as a hybrid-multi-cloud service. The design of this service segregates data based on data sovereignty and privacy concerns.

Individual application workloads are spread across multiple hypervisor clusters. What the end user sees as a single application thus has a minimum of 10 workloads spread across an unknown number of hypervisor clusters, with potentially dozens of individual IT components able to impact application performance.

Identifying application performance problems requires our systems administrator to have a deep understanding of the supporting infrastructures and services used for the application workloads. If components of this application live on multiple infrastructures, this could mean that our systems administrator has an in-depth understanding of all the individual on-premises infrastructure components upon which workloads for the application might run, as well as any public cloud infrastructures that might host workloads for the application.

That's a lot for any one person to know, especially considering that our systems administrator must also know the application itself, inside and out, forward and backward. The traditional solution to this problem has been to hire IT teams with multiple infrastructure specialists, in addition to the application administrators.

Unfortunately, it's no longer economically feasible for most organizations to retain specialists for all infrastructure silos. This is especially true where multiple (e.g., public cloud and on-premises) infrastructures are in use, and today's organizations use many different types of infrastructure from a large number of vendors.

Making matters worse, it's increasingly rare for a IT team member to be assigned to only one application. Having an entire IT team devoted to a single application is almost unheard of, being reserved for organizations like Netflix or Facebook, where the "application" in question is functionally the entire business, and the application that the customer sees isn't a mere 10 workloads, but hundreds of thousands of workloads.

The average enterprise systems professional working in the average organization must keep an ever-growing list of applications functioning, and do so economically. The number of workloads under management is constantly growing in most organizations, while the number of IT team members isn't. Thus, an enterprise IT team needs the right IT visibility tools to do their jobs, and this includes monitoring tools capable of analyzing application outages and performance problems from the application workload perspective, all the way down to how different hypervisor storage arrays respond to different network conditions.

The average enterprise systems professional working in the average organization must keep an ever-growing list of applications functioning, and do so economically.

TRADITIONAL MONITORING TOOLS DON'T CUT IT

Just as it's rare for organizations to dedicate even a single IT systems professional to an application, it's also rare for systems administrators to keep applications running with only the built-in tools provided by the IT infrastructure supporting their application. Enterprise IT has access to a wide range of monitoring tools, security information and event management (SIEM) platforms, and more. Without these tools, it would be impossible for most organizations to sustain the current ratio of workloads per IT professional.

Unfortunately, most organizations are at their limits with their current enterprise systems toolchains. Most organizations are still using tools that aren't up to the challenge of monitoring, much less managing today's applications.

Traditional infrastructure monitoring has no application awareness, while traditional Application Performance Monitoring (APM) cannot see below the host/VMs. In most organizations, only 10-20 percent of applications are monitored with APMs, primarily due to the cost of application monitoring products.

Where application monitoring is deployed, it often suffers from the opposite problem of infrastructure monitoring solutions: application monitoring traditionally takes an end-user-centric approach. It focuses on how the application performs for the users; it doesn't have much, if any, insight into the underlying infrastructure.

Making matters worse, traditional infrastructure and application monitoring tools are complicated, requiring significant expertise to use. While these tools are useful for surfacing relevant data for systems admins, the siloed nature of traditional monitoring leaves it up to the admins to correlate data from multiple sources and identify a root cause.

Today's organizations already have enough problems finding IT experts; requiring experts (or at least expertise) in particular infrastructures and application monitoring tools isn't helpful. Enterpise IT teams need modern tools and analytics that provide application-centric infrastructure management capabilities.

Most organizations are still using tools that aren't up to the challenge of monitoring, much less managing today's applications.

VirtualWisdom Understands

It's easy to write about what should exist; it's much harder to talk about the real-world difficulties of IT. It's unlikely that any monitoring or management tool will ever provide a perfect mix of infrastructure and application awareness. There are simply too many variables.

It's economically feasible for vendors to create application monitoring tools for extremely popular applications with hundreds of millions of deployments, or tools for outrageously expensive applications where even costly monitoring would be a fraction of the application's price. For some vendors, it's economically feasible to build in support for several popular applications. It will, however, never be possible to build application awareness for more than a fraction of the applications deployed.

The same is true for infrastructure monitoring. While it's entirely possible for vendors of monitoring applications to build in support for the most popular infrastructure components, the number of complete infrastructures (let alone infrastructure components) is increasing faster than the number of monitoring products that vendors can produce. What's needed is a modern approach to realize the visibility and control that modern data centers require.

What's needed is a modern approach to realize the visibility and control that modern data centers require.

GETTING A HANDLE ON INFRASTRUCTURE

VirtualWisdom is a modern app-centric IT management platform that provides complete, end-to-end hybrid IT infrastructure control and visibility. VirtualWisdom discovers, identifies, combines, and correlates workload traffic and event data, using network wire data, server event data, and infrastructure performance data.

Collecting data from all infrastructure elements (e.g., applications, compute, memory, network, and storage), VirtualWisdom has been developed with an in-depth understanding of how complex systems interoperate and how application service resources and infrastructure service resources are dependent upon each other.

All data collected by and about application and infrastructure components is fed through the VirtualWisdom WorkloadAI platform, which allows VirtualWisdom to apply analytics to solve problems at scales that would be otherwise impossible. The VirtualWisdom WorkloadAI platform can analyze data from multiple infrastructures in real time, providing insights into applications and services with workloads spread across multiple public clouds and on-premises data centers.

VirtualWisdom's application awareness affects its understanding of the underlying infrastructure. For example, VirtualWisdom can understand that multiple workloads combine to form a service, and that multiple services can combine into what appears as a single application to an end user.

VirtualWisdom can understand not only which infrastructure components support which individual workloads, but also how these workloads interact to form a service, how those workloads

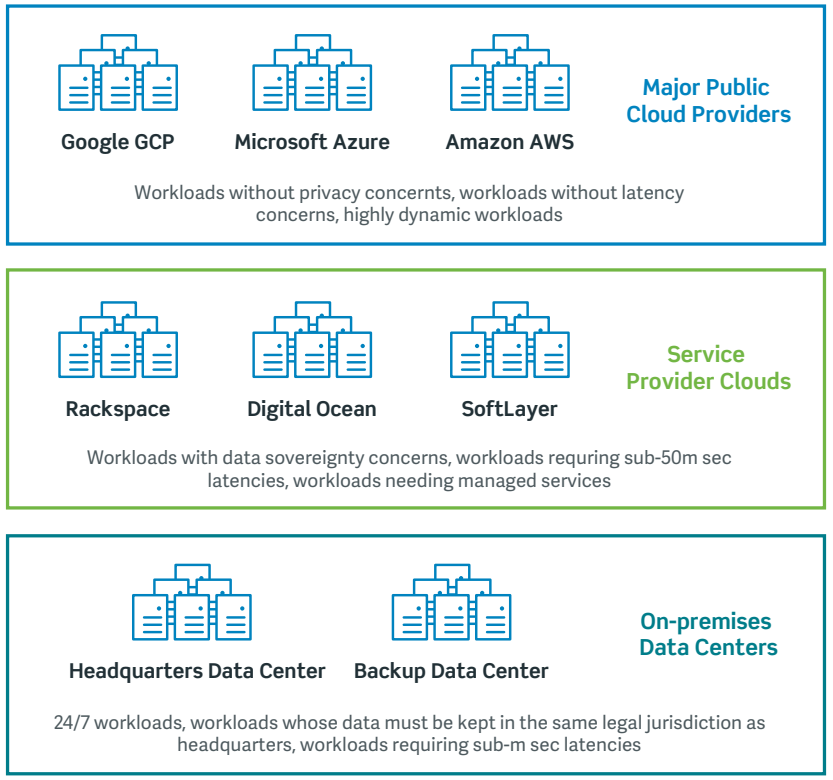

Figure 3: Different workloads have different requirements.

interact with other workloads on shared infrastructure, the relative importance of individual workloads to a service, and the relative importance of services or applications to one another.

VirtualWisdom uses multiple data sources, including those provided by individual infrastructure components, such as VMware vCenter, IP networks, network fabric switches, and individual storage arrays. VirtualWisdom also ensures the performance of

applications, allowing single tiered alarms to be shared across every infrastructure service supporting that application. This allows for a "set it once and forget it" approach.

VirtualWisdom collects, correlates, and analyzes this data in real time, offering a complete understanding of IT infrastructure spanning public clouds and on-premises data centers, all in a single pane of glass. The VirtualWisdom WorkloadAI platform uses this data to make predictions about ongoing performance and to inform IT pros of the risk/reward of proposed changes.

GETTING A HANDLE ON APPLICATIONS

VirtualWisdom understands applications. The VirtualWisdom platform has integrations into leading application management platforms, including AppDynamics, Dynatrace and others, and can identify how applications use infrastructure. Performance wire data analysis and baselining are both critical to characterize unknown applications and offer IT teams insight into entire categories of applications that monitoring solutions would otherwise completely ignore.

Perhaps more critical than the ability to isolate and monitor unknown applications is VirtualWisdom's understanding that applications on shared infrastructures have mutual performance impacts. Understanding the interactions between applications, both within and outside the boundaries of shared infrastructure, is an essential feature of the VirtualWisdom WorkloadAI platform.

By combining infrastructure awareness with application awareness, VirtualWisdom can help IT teams identify how a change in one part of an infrastructure can affect multiple applications, helping to prevent performance impacts and outages *before* they occur. VirtualWisdom also performs application and infrastructure discovery and mapping. Application discovery methods include SSH, WMI, and ServiceNow CMDB, as well as pulled information from vSphere distributed switches, NetFlow data, AppDynamics, and Dynatrace.

VirtualWisdom uses the information gathered both from applications and infrastructure to offer more than just monitoring and alerting. VirtualWisdom provides IT teams with predictive capacity management, letting them know when existing infrastructure is reaching its limit and when additional capacity must be brought online to meet requirements.

VirtualWisdom offers workload placement recommendations and even provides recommendations for storage port balancing and VM balancing, helping organizations achieve a 30% - 50% savings on infrastructure costs.

By combining infrastructure awareness with application awareness, VirtualWisdom can help IT teams identify how a change in one part of an infrastructure can affect multiple applications, helping to prevent performance impacts and outages *before* they occur.

THE HUMAN FACTOR

The real challenge for IT monitoring and management applications today isn't collecting data, or even processing all the information collected. The real challenge is making sense of all this information so a human being can understand and act on it. Remember also that not every person will be interested in the same information.

With VirtualWisdom, insights emerge through highly customizable dashboards. Each dashboard contains content and visualizations that are optimized for different audiences.

Infrastructure operations teams, for example, might focus on utilization, capacity planning, and infrastructure health. The break/fix concerns of infrastructure teams range from items like failed hard drives and network ports to the failure of compute nodes, firewalls, and internet connectivity.

In organizations where administrators are assigned by area of specialization rather than by particular applications, infrastructure operations teams frequently serve as the hub of all IT activity. Beyond capacity planning and break/fix items, they may also investigate performance metrics and root cause analysis that's just granular enough to know which other IT team members to pull in on a problem.

Storage professionals are one such specialty that IT operations teams may seek. Larger organizations likely have in-house storage administrators, architects, and technical engineers, but smaller organizations are more likely to have contracted this role to a managed service provider. Storage admins will be looking for dashboards that quickly convey storage performance, current capacity balance, calculated optimal capacity balance, storage system health, and availability of both storage and network resources to various workloads.

Senior IT team members, such as data center architects, may also be required. Senior team members are typically asked to solve the hard problems, like performance issues, rather than straightforward outages. Diving deep into a problem means that architects are likely to need granular wire data and machine data, full-stack visibility, and problem resolution tools. They want to take proactive steps to optimize the infrastructure for cost and performance.

In contrast, management rarely wants the nitty-gritty details. A dashboard that meets management needs is likely to include end-to-end visibility and complete infrastructure health, total infrastructure utilization, recommendations regarding infrastructure expansions, and any cost controls, should those be applicable to the current infrastructure.

Each organization coordinates its IT teams differently. To accommodate this, VirtualWisdom has dozens of dashboard templates out of the box, but is completely customizable to meet the needs of various stakeholders within an organization.

> The real challenge is making sense of all this information so a human being can understand it.

VIRTUALWISDOM UNDERSTANDS

Presenting someone with tons of raw data and then asking them to solve a problem based on that data isn't making life easy for that person. With the right processes and experience, the average person can find correlations in datasets and eventually identify the root cause of a problem. Unfortunately, humans are notoriously slow at this, which is one of the main reasons for inventing computers in the first place.

Expertise allows a person to more accurately filter out data that's irrelevant to a given problem and to more readily recognize data points that are likely to be relevant. This ultimately reduces the number of necessary passes over the data, but it doesn't change the speed at which that person can process data. Unfortunately, no matter how expert a person is, there are hard limits to how fast they can process data. Computers can sift through data volumes that are functionally impossible for manual human labor. What computers have traditionally been missing is expertise and context for the problem they're addressing; filtering signal from noise requires a great deal of understanding.

Learn more by visiting
www.virtana.com

VirtualWisdom is a virtual infrastructure and application performance monitoring expert , one capable of handling data no human could realistically sift, doing so fast enough not only to help resolve IT outages and performance problems, but also to proactively help prevent them in the first place.

VirtualWisdom understands the physical infrastructure upon which applications run. VirtualWisdom also understands the applications that run on that infrastructure, either because of native integration/discovery or because it can baseline that application's footprint across compute, memory, storage, and network usage through other integration means.

Most importantly, VirtualWisdom understands how to present information to humans. The end result is a product that provides application service assurance, predictive capacity management, workload and capacity optimization, as well as problem resolution and avoidance.