



Spotlight

Spotlight Paper by Bloor

Author **David Norfolk and Paul Bevan**

Publish date **January 2018**

Instrumenting the Virtualised Datacentre ...for Performance SLAs

“

**However fluffy
the cloud, somewhere
real data is moving
off real disk drives,
down real wires,
to real computers/
switches in noisy,
air-conditioned,
very non-fluffy,
concrete buildings
and, here,
performance
matters.**

”

Executive summary

This document is addressed to Infrastructure Managers (looking after mission-critical applications) responsible for monitoring SLAs (Service Level Agreements) for virtualised infrastructure; their staff; and (probably executive summary only) associated business managers.

The nub of the business issue we spotlight is that performance affects infrastructure risk just as much as availability (probably more so) but vendors of virtualised infrastructure platforms and services (and even an in-house technology provider) are loath to provide effective performance SLAs. In fact, even if they do provide performance information, can it be relied on to be unbiased? And, is it siloed rather than holistic? In the increasingly virtualised world, in order to deliver services reliably, companies need seamless access to performance metrics across their entire infrastructure, whether on-premises or in private/hybrid/public cloud.

We recognise that many organisations are still using on-premises legacy infrastructure; but we believe that the trend is firmly towards virtualised infrastructure and Cloud-enabled interfaces, even if a business case for Cloud cannot be made at present, in a particular situation. In any case, the availability of robust performance metrics is important, even for legacy infrastructure.

An essential enabler for the journey towards Cloud is Big Software, which means the use of software-defined, virtualised infrastructure. Or, in other words, general-purpose computers configured to act as different hardware devices by the software programs running on them. Deployment of tools can be programmed, the metrics they deliver can be programmed, and the metrics themselves can be measured (usage metrics are key; if a metric isn't being used to benefit the business, it is a wasteful overhead).

Big Software is inclusive – if a less sophisticated start-up, say, is thinking in terms of appliances, Big Software lets you place a software-defined tool onto

bare metal – and you have a software-defined appliance which can evolve into something more sophisticated if the business needs it. Big Software also promotes end-to-end traceability, so whole-system dashboards can present metrics and alerts in purely business terms and still support drill down to, say, packets on a fibre-channel link for the technicians in the organisation.

Performance starts with data moving between storage (whether virtualised or not) and processing. If that is too slow, performance tuning higher up the stack can't help much. What is needed then, in part, is technology- and vendor-agnostic performance monitoring, ideally at the ethernet and fibre-channel fabric level, so no vendor can withhold information. The tool (or integrated tools) providing such monitoring should, ideally:

- Place no load on the system – so it should be agent-less and “*out of band*” (an agent is a piece of software running in the background on a device, collecting management information; out of band communications do not share bandwidth with the business workload);
- Operate in real-time, with the lowest possible latency;
- Be highly secure, so it should operate on packet headers only, not the payload and offer tested/certified error-proof security;
- Be usable by all stakeholders, so with role-based dashboard access and a drill-down interface (from traffic light to engineering detail);
- Provide exception reporting in terms of variations from a “*normal*” workload pattern;
- Document and learn from the fixing of problems (Machine learning);
- Offer “*root cause analysis*”, and avoid error message storms;
- Be part of an active ecosystem/partnership program, which reduces risk for its customers and ensures holistic integration with other monitoring tools.



Companies need seamless access to performance metrics across their entire infrastructure, whether on-premises or in private/hybrid/public cloud.





Performance metrics should be aggregated and prioritised in business-level dashboards, with “traffic light” service indicators... role-based access... and drill-down.



Such low-level infrastructure monitoring, though necessary, is not sufficient. Agents still have a place, perhaps addressing specific problems. Once you've addressed storage issues and bottlenecks, other issues and bottlenecks can appear elsewhere, so it is important that companies take a holistic approach to monitoring performance SLAs throughout the stack. So, data infrastructure monitoring tools, which deal with the issue this Spotlight focuses on, must have APIs for integration with other, more specialist, software-defined tools that allow a holistic management of performance.

It is also important that performance metrics, even from the low-level infrastructure, can be related back to real business issues. There is little point in devoting resources to fixing a potential infrastructure bottleneck which isn't actually impacting the business, if the customer service is being immediately impacted by issues further up the stack. In other words, it is hard for the business to prioritise issues if they are only described in technical terms, and (usually) the business should be prioritising issues in terms of business impact. In practice, this means that performance metrics should be aggregated and prioritised in business-level dashboards, with “traffic light” service indicators (OK-warning-serious); with role-based access throughout the company; and with drill-down to technical metrics. This is sometimes called “*Application-centric Infrastructure Performance Management*” (IPM).

These days, you can find vendor-agnostic tools to deliver the required performance metrics and SLAs, but you generally have to look for them. They are often not available by default on a platform.

Infrastructure performance SLAs, an emerging issue

Companies are moving towards public, hybrid and private cloud, as they have to become increasingly *Mutable* (change-embracing). Even if they decide to remain with on-premises solutions, often for regulatory reasons, for now, they will increasingly build for the Cloud, so they have the freedom to move onto Cloud solutions if, as or when this makes sense. Cloud is not really an end in itself, but it is the freedom to run workloads wherever it makes business sense to run them. And, of course, if you can't manage service performance effectively across existing in-house silos, how will you ever manage service performance SLAs in the cloud?

Usually, these days, this means that Hybrid Cloud is the target, and Bloor sees this continuing for the foreseeable future. The emerging issue is that delivering services to customers depends, ultimately, on the performance of the storage infrastructure, although it is then impacted by performance higher up the stack too, of course. However fluffy the cloud, somewhere real data is moving off real disk drives, down real wires, to real computers/switches in noisy, air-conditioned, very non-fluffy, concrete buildings. And, as Cloud Spectator points out here [<file:///E:/Documents/Copy/Bloor/VirtualInstruments/2017-Cloud-Spectator-EU-Report-05-09-17.pdf>], "Public cloud service providers (CSP) purport to offer instantaneous, scalable virtual infrastructure with utility billing. In reality there is wide variance in cloud performance".

So, in order to manage customer-focused service levels, a company has to manage infrastructure service levels, and not just manage availability but performance – even on virtualised infrastructure that may not be directly under its control. If it uses Cloud infrastructure, it then hits a problem: cloud vendors will happily provide availability SLAs (of varying quality, but that's not the issue here) but

won't provide performance SLAs. Nevertheless, performance issues are common and more likely, in practice, to drive customers to the opposition than availability issues, which are, given reasonable luck and competence, fairly rare.

In addition, there is little standardisation in the semantics of Cloud (and on-premises) performance claims. Most obviously, is "high availability" less than 5 minutes downtime a year or does "planned downtime" (an hour or so a week installing updates) not count? The net result is often expensive over-provisioning, just in case the cloud elasticity or other characteristic claimed by a particular CSP doesn't eventuate in practice.

This issue remains even if your storage infrastructure is in-house, despite the illusion of manageability this provides. If you look at on-premises infrastructure as a service to the business, and why wouldn't you, what SLAs does it give you? Is there any formal contract with the on-premises infrastructure facility as to the performance it will provide and, if performance falls below "acceptable", is time wasted in finger-pointing around all the groups that may be at fault? Managing infrastructure may even be easier with Cloud, paradoxically, because even though Cloud infrastructure may

be outsourced, Cloud jurisdictions, SLAs, and contracts should all be well-defined. Even a missing SLA provides a point of certainty compared to a fuzzy, ill-defined and unenforced one such as you often find, in practice, for on-premises infrastructure.



However fluffy the cloud, somewhere real data is moving off real disk drives, down real wires, to real computers/switches in noisy, air-conditioned, very non-fluffy, concrete buildings.



At a technical level, in order to claim that its infrastructure is well-managed, a company must take ownership of the performance metadata it needs to manage infrastructure performance. It needs to be able to check that infrastructure performance SLAs, if it can get any, for relational databases, block-file, object storage and so on are adequate and are being met. It should also be able to identify infrastructure performance issues itself, if necessary, and address them efficiently and, ideally, proactively.

At a business level, in order for a customer to trust a company to deliver a service that is appropriately performant for its needs, that trust must extend right through the stack from the application to the infrastructure. Trust comes with good management; and good management needs good metrics from effective monitoring of availability, costs and performance. Unfortunately, vendor-neutral, trustworthy, performance metrics and SLAs across the whole business are not always that easy to obtain, especially at the infrastructure level and with increasingly virtualised infrastructure.

As W. Edwards Deming said: *"In God we trust; all others must bring data."*



In God we trust; all others must bring data.



W. Edwards Deming



Why should you care about infrastructure performance?

In the end, the only real reason to care about infrastructure performance is that it has a significant impact on the business, which means that even low-level performance metrics should be accessible through drill-down from business-level performance dashboards (or equivalent). Business impacts could include, for example:

- Financial traders “losing” millions of pounds if their trading systems slow by even a millisecond.
- Websites relying on advertising based on “eye-balls” and “time on site” losing revenues if viewers click away because of slow response times.
- On-line retailers losing customers and revenues if page load times are slow.
- End-users expecting “good response times” irrespective of whether they are working on an internal admin/operation system or an external customer facing one, and then, if they don’t get it, losing morale and performing badly for the business.

However, it’s not just infrastructure, of course. Performance issues can arise in the application, because of poor design and lack of scalability (owing to issues with threading, disk i/o, etc.). They can arise in the switches and routers (generally through software configuration issues, but also because of the lack of proper capacity). They can arise in the servers (through configuration issues around virtualisation and prioritisation). And, a particular focus of this Spotlight, they can also arise in the in the storage sub-system through:

- Poor paths to disk.
- Variation in performance across data types because different storage types (relational database, block-file, object storage) perform very differently with heterogeneous workloads.
- “Noisy neighbours”, where you are impacted by poorly designed applications sharing your bandwidth or other resource contention issues.

Although Cloud Spectator (op. cit.) claims that many of the largest CSPs address this, some resort to performance throttling in order to deliver a consistent user experience (which is important), and this implies that the ultimate performance capability of the infrastructure isn’t available to users.

Once you start addressing these issues, you need to care about problems with siloed operational teams and vendor or technology specific tools. The business does not want to hear “not my problem, try that lot over there” or “problem not caused here, go away”. It wants the business issue fixed, quickly and effectively. It wants the root cause of an issue identified quickly and collaboratively, and to see someone put onto addressing it. It doesn’t want to see developers and technicians playing the “it’s not my fault” tag game.

An alternative payments company doesn’t worry about one transaction being slow, in itself. It worries that once a customer uses a credit card because the alternative payment is slow, they stick with using credit cards – and perhaps influence their friends to do the same. In such circumstances, the company won’t be concerned about finding someone – or a department – to blame, it will want to inculcate a network of **TRUST** through the stack from customer to infrastructure provider – which ensures that people are collaborating on solving the business issue, not on blame avoidance; and that people are not satisfied if just their little bit of the chain works well and so will share information and metrics along the whole chain.

Infrastructure has to be managed by someone, and this management must be based on good-quality, vendor-independent, end-to end metrics. In practice, this is hard enough on-premises, especially with virtualised infrastructure; but Cloud performance is not a given and it is something that you have to manage, even if you are using a cloud service provider.



The business does not want to hear “not my problem, try that lot over there” or “problem not caused here, go away”. It wants the business issue fixed, quickly and effectively.



What we can do about IPM



It is important that an IPM (Infrastructure Performance Management) platform or tool suite offers a low cost of entry and easily manageable options for small start-ups and departmental initiatives, which can “grow by success”.



Obviously, there are a whole range of performance management and reporting tools that can contribute to solving this business issue and we can't talk about them all here. Nevertheless, they will all need to have some common characteristics:

- They need to be able to identify where the problem seen by the business is happening in real-time – or be able to talk to a specialised, low-level Application Performance Management or Network Performance Management tool that can.
- They must have well-defined, well documented APIs, for tool integration, even if they do a good job of reporting technical performance issues in business terms themselves.
- They must operate across technology and organisational boundaries and be vendor-independent, if necessary. They should be able to use, for example, monitoring information from a cloud service supplier, but they should also be able to collect this at a more granular, vendor-independent level, perhaps by monitoring actual packet headers.
- They must not impose a significant overhead on the system. This usually requires that an agent-less option must be available, with out-of-band reporting. This means that management communication of performance issues doesn't use bandwidth that is being used for business communication. A hardware or appliance implementation should be available, although agents, in-band alerts and software implementations all have a place too. Customer choice and flexibility is key.
- With Cloud services, in particular, even if enterprises can't see what is going on inside their cloud service provider (CSP), they should be able to assure themselves about the tools the CSP is using to get that real-time visibility itself. They should also own the performance metadata, as well as the data, associated with their use of the service.

The point is that we need end to end solutions, across technology boundaries. These need to be firmly rooted in business impact, user experience management, and business terminology.

This can be referred to as “*Application-centric Infrastructure Performance Management*” (IPM). This needs to be both top-down and bottom-up. So, you need to understand the impact of

storage technology decisions on business application performance; but you also need to understand business workloads, which ones are most important to the business, and which ones stress the infrastructure the most. At the highest level, performance, as perceived by the business, comes from matching the most important workloads to the most appropriate infrastructure.

Typically, an organisation will need role-based access to dashboards supporting impact analysis and offering drill down from “customer widget processing service impacted” to an issue with, say, packets being transmitted on fibre-channel.

It is important that you think about what you want to achieve with IPM (Infrastructure Performance Management) before you start. Select tools which match your requirements, not just those that the vendors want to sell you. Nevertheless, you need to look beyond your status quo – allegedly Steve Jobs succeeded by giving his customers what they didn’t know they wanted yet – although it is also important that your “10-year future vision” doesn’t distract you from keeping your paying customers happy with the service you are giving them today. Seek independent advice about what is achievable today and will be possible in the future. By all means listen to what the vendors say (if it isn’t insightful and worth listening to, change your vendors), but an organisation must take ownership of its own vision, not abdicate responsibility to a plausible vendor.

Then you can implement your own application-centric IPM solution across the whole stack, delivering immediate benefit to users today as well as future capabilities. At Bloor, we think that it is important that an IPM platform or tool suite offers a low cost of entry and easily manageable options for small start-ups and departmental initiatives, which can “grow by success”. Continuity is important. Implementing a small, but cheap and manageable, solution which doesn’t scale, then ripping it out and replacing it by something large and expensive, which is hard to manage, is not particularly effective. Besides, developing a small, manageable, entry-level solution encourages a vendor to keep its large, enterprise-capable solutions manageable and user-friendly.

It is then vital that you review your implementation against your original vision, rather than just what your technology does well. Baseline your systems. Identify what worked and what

didn’t – and decide what you will do about what didn’t work.

In a 21st century “mutable organisation” (an organisation in a constant state of evolutionary change and reinvention as its business environment changes) the environment supporting this is Big Software. Increasingly, infrastructure hardware is implemented as virtualised general-purpose computers running specialising software, which can be deployed and configured by running software programs. So, a hardware box can be a “bare metal” component running software that makes it look like a load balancer or performance monitor; but it may equally well be a virtualised software emulation of the “bare metal” running exactly the same software. It is, however, important that all this software utilises “no code” or “model driven” visual programming approaches – an organisation doesn’t want to replace inflexible, unmanageable physical hardware with unmanageable, inflexible Big Software spaghetti.

Big Software or “software defined” is a generalised term for all that. It is, in part, about linking all of the tools needed for IPM together in a virtualised, software-defined, holistic management system, using well-defined APIs (Application Programming Interfaces). So, a software-defined load balancer, for example, can be deployed in interesting and innovative ways by a sophisticated service provider, because its capabilities and configuration can be changed quickly and easily using software. However, the same load balancer can be implemented on bare metal and deployed as a simple appliance by less sophisticated organisations.

Finally, it is important that the vision of software-defined metrics and complete transparency across the stack, end-to-end is also managed. Usage metrics are vital. The metrics and information gathered must be USED – to validate SLAs, to highlight impending issues, and to facilitate 2-way TRUST between all of the stakeholders in the organisation. Big Software capabilities that aren’t being used by the organisation represent a waste of resources.

“
Big Software capabilities that aren’t being used by the organisation represent a waste of resources.
”



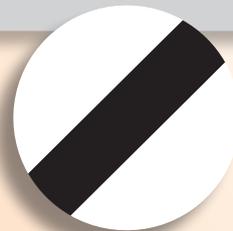
Other issues

It is possible that major cloud and infrastructure vendors will purchase (or re-invent) the sort of tools we are talking about here and provide performance SLAs, together with the data needed to validate their achievement. Even if this happens, there will still be a need for vendor-independent validation of the achievement of these SLAs – you can outsource execution of infrastructure performance management, but you can't outsource responsibility for its governance.

More interesting, perhaps, is the possibility of the rise of trusted Cloud Service Integration Brokers, that will select Cloud providers and translate whatever SLAs they provide into the actual availability and performance SLAs needed by different business models. Such service brokers will probably use the tools we talk about here, or recommend tool purchases to their clients, in order to manage the SLAs for their clients and will also provide performance data to their clients, with their own SLAs.



Interesting, perhaps, is the possibility of the rise of trusted Cloud Service Integration Brokers.



Customer use cases

It became clear while researching the industry for this Spotlight that there were a large number of vendors who have sophisticated tools that focus on specific areas such as network performance management, application performance management, server management etc. But there were very few who had solutions that provided a holistic view of multiple infrastructure components. This was borne out by comments from the Chief Innovation Officer at EdgeConnex, Phil Lawson-Shanks, who confirmed that when it was looking for such holistic capability back in 2015, it found it necessary to make the decision to develop its own infrastructure performance management capability in-house.

However, we can now see that some organisations have started to implement IPM as illustrated by Use Cases we have seen from, for example, Virtual Instruments with its VirtualWisdom application-centric IPM solution (this is one of the few vendors we have come across with capabilities that approach the requirements we have described in this document). It is important to note that such holistic IPM tools do not necessarily negate the requirement for deep dive specialist tools like, for example, Application Performance Management or Network Performance Management.

De-risking technology growth

Primary Actor: Leading ecommerce payment business that allows payments and money transfers to be made securely through the internet. The payment service is as an electronic alternative to traditional paper methods such as cheques and money orders.

Scope: Online payments system operating in over 100 markets with millions of accounts.

Level: Business credibility

Story: While this company is one of the leaders in this space, consumers have many choices when paying online. If one online payment option is slow or has an outage, consumers will move over to another option, such as paying by another credit card – and may never come back.

This results in a technology challenge for IT, which must support growth and ensure that new applications are rolled out quickly and reliably, that they are up and available to users 24x7x365, and that application response times to the users meet their expectations. Fundamentally, the IT administrators need transaction visibility through the I/O system,

getting meaningful data about their infrastructure on a continuous, real-time basis. Without this data and measurement system, application performance is difficult to optimise, and every application demands the most expensive, tier 1 storage.

This challenge was addressed by installing an instrumentation, monitoring, and analysis platform which can holistically assess the entire multi-vendor SAN (Storage Area Network) infrastructure, and get the data necessary to make intelligent decisions about capacity, utilization, and performance for every layer of the infrastructure – network, server, storage, and applications. The platform gets its fundamental data from the network switches via SNMP (Simple Network Management Protocol), and from the physical layer via an optical splitter (which provides real-time, bit-for-bit copy of all traffic), thus imposing no load on the business processing. Every Fibre Channel frame (transaction) can be analysed in real-time to provide dozens of critical statistics about storage traffic, transmission errors, and SAN read and write latency.

Customer benefits achieved included:

- Virtual elimination of wasteful “finger pointing” when addressing issues.
- Real time tracking of SAN performance as changes are introduced.
- Proactive addressing of emerging problems (often before the customer notices anything wrong)
- For a datacentre migration, benchmarking of the application performance at the old datacentre, tracking of latency metrics during the migration, and assurance of no performance degradation once the migration/consolidation was complete.
- Validation of vendor claims and elimination of compatibility issues in a test lab before deploying new technology/services.
- Reduction in the need for wasteful over-provisioning, while giving assurance that customer SLAs will be met.
- Facilitating a move to a cloud or chargeback model, to provide a QoS (Quality of Service) model for customers.

Customer use cases ctd...

A productive partnership Between IT and the Business

Primary Actor: A large retail organisation.

Scope: Trust and transparency across an organisational structure.

Level: Company strategy

Story: The new Head of Enterprise Computing for this company found, on his arrival, *“an infrastructure that was ageing, highly bespoke, with multiple vendors and lots of islands of infrastructure supporting many business units. The risk position was high in respect of culture and capability. Processes were not correctly established or properly documented; there were few standard delivery routines, and insufficient disaster-recovery testing.”*

A ground-up infrastructure review identified storage issues as having a major business impact and, therefore, legacy storage was the first priority for upgrade. Nevertheless, he saw that the lack of rigorous management practices around storage health, utilization, and performance, the use of ad-hoc fixes, and generally poor visibility of the storage infrastructure, was actually a key contributor to the storage-related issues he was finding, as much as the age of the hardware. He was already familiar with a Service Performance Management tool suit from previous positions and introduced these tools to support three important new management practices:

1. Working face-to-face with the business to develop a detailed analysis of its apps and supporting infrastructure – for example, documenting which servers support which applications.
2. Establishing the baseline application performance on the legacy storage arrays, before they were optimised or replaced. This enables continual managed improvement of services – if you don't have a baseline, how do you know whether the work you are doing is improving anything?

3. Holding monthly application performance reviews with the business units, which helps to educate them in how their infrastructure is performing and in the impact of their applications on the infrastructure.

The Service Performance Management platform (and the tools supporting it) are now being used widely; and, increasingly, proactively, in the organisation. This is largely because the platform is open to the business and allows it to gain detailed insights into the application infrastructure and how well it is working.

Customer benefits achieved:

- Cost reduction was not the driving factor here. Rather, by creating best practices and by putting the right foundations in place, the Enterprise Computing group is starting to gain the trust of the business. This allows better management of, and outcomes from, its infrastructure.
- Improved and more proactive capacity planning, instead of ad-hoc over-provisioning, is becoming the norm.
- A key strategic benefit of making the infrastructure performance transparent, both to the IT staff and the business units, is that there is a common foundation for collaboration and planning, in a climate of trust, based on objective and accurate performance data.

Supporting SaaS Operations

Primary Actor: Large North America Cloud-based business software organisation, with over 10,000 employees and more than \$20billion yearly revenue

Scope: IT infrastructure managers responsible for reliable, always-on, telecommunications Software as a Service (SaaS), and Infrastructure as a Service (IaaS), delivery.

Level: Strategic operations architecture.

Story: With its previous tools, this SaaS provider had suffered from a lack of comprehensive visibility into its channel fabric. It implemented a new toolset, with multiple data collection sources, to fill the gap in fabric visibility. This reported on the overall health, utilisation and performance of the organisation's critical infrastructure.

Its new tools helped this SaaS provider to:

- Identify and address several areas of concern around: high utilisation; host multi-path imbalance; and, transmission errors (at the host, switch & storage port level). These issues had serious potential for impacting the availability and performance of mission critical systems.
- Report on all aspects of the flow control equation: for IO size; queue depth settings; and, buffer credit depletion. Data at this detailed level is critically important as the provider upgrades device speeds from 4 to 16 GB/sec and adds faster (*“all flash”*) storage devices.

- Give the SAN team strict SLAs with IT operations. These require impartial, “wire level” (that is, low level, detailed), data, showing end-to-end exchange completion times with a breakdown of the latency within the infrastructure, from initiator to target to logical storage device (or LUN, the unique identifier given to the separate devices, or logical units, accessed with the Fibre Channel protocol).

Customer benefits achieved:

- OpEx benefits, achieved through comprehensive visibility into the overall performance of the infrastructure, eliminating reliance on lower-fidelity “performance” data, coming from device specific tools and, often conflicting, data from disparate Application, Database and Storage teams.
- CapEx benefits, from the availability of impartial performance data, which allows the organisation to “right size” its infrastructure in response to the changing demands of its application workloads.

Reducing Storage Migration Risks

Primary Actor: UK-based, vertically-integrated supermarket chain that has been in business, changing and adapting to meet its customers’ needs, for over 100 years.

Scope: Storage Infrastructure

Level: Technical IT infrastructure

Story: In order to provide performance verification capabilities, its IT team created performance dashboards for a number of this supermarket’s business-critical applications. These facilitate the creation of base-lined performance benchmarks as the infrastructure evolves, both pre- and post-migration. This means that performance or configuration issues can be identified proactively, before business services are impacted. Overall, this supermarket chain’s tools give it objective evidence that its evolving IT infrastructure continually meets or exceeds its performance objectives and “non-functional requirements.”

It can now plan for additional growth and capacity based on definitive performance information, rather than on raw volume numbers and guesswork. Its head of storage says, *“There’s a strong temptation to simply throw hardware at a performance issue, but often that just moves the bottleneck rather than solving the problem. [Now] we can make accurate and strategic investments based on sound performance information.”*

Customer benefits achieved:

- Rapid and accurate identification of significant performance issues during testing, so problems can be corrected before the production migration – and before they impact the business.
- Reduced migration risk from monitoring of performance in real time as migration proceeds.
- Reduced CapEx; that is, cost avoidance through low-risk, proactive, capacity planning guided by actual performance requirements, instead of high-risk over-provisioning, based on worst-case guesses about expected volumes.
- Successful migrations of business-critical applications to new storage environments.
- Optimisation of the performance of these new storage environments.
- Reduced risk of change in the new environment, post-migration, through constant monitoring of real-time SAN performance during future migrations, upgrades and reconfigurations. These ensure that changes do not impact application performance, as experienced by business users.
- Comprehensive visibility, for managers, into system-wide performance.

Summary

Increasingly, Bloor sees application-centric Infrastructure Performance Management as a major business issue facing emerging *Mutable* (change-embracing) organisations. This is because the business can be seriously impacted by even low-level infrastructure issues and these are increasingly hard to manage as technology becomes virtualised and the use of hybrid cloud services increases.

The solution of this issue lies in Big Software – software-defined infrastructure with vendor-independent performance monitoring tools, operating across the entire stack, with role-based access from management through to technician.

FURTHER INFORMATION

Further information about this subject is available from www.bloorresearch.com/update/2354



About the authors

PAUL BEVAN

Research Director: IT Infrastructure

Paul has had a 40-year career in industry that started in logistics with a variety of operational management roles. For the last 33 years he has worked in the IT industry, mostly in sales and marketing, covering everything from mainframes to personal computers, development tools to specific industry applications, IT services and outsourcing. In the last few years he has been a keen commentator and analyst of the data centre and cloud world. He is also a non-executive director in an NHS Clinical Commissioning Group.

Paul has a deep knowledge and understanding about the IT services market and is particularly interested in the impact of Cloud, Software Defined

infrastructure, OpenStack, the Open Compute Project and new data centre models on both business users and IT vendors. His mix of business and IT experience, allied to a passionate belief in customer focus and “grown-up” marketing, has given him a particular capability in understanding and articulating the business benefits of technology. This enables him to advise businesses on the impact and benefits of particular technologies and services, and to help IT vendors position and promote their offerings more effectively.



About the authors

DAVID NORFOLK

Practice Leader: Development and Governance

David Norfolk was working in the Research School of Chemistry at the Australian National University in the 1970s, when he discovered that computers could deliver misleading answers, even when programmed by very clever people. His ongoing interest in getting computers to deliver useful automation culminated in his joining Bloor in 2007 and taking on the development brief.

Development here refers to developing automated business outcomes, not just coding. It also covers the processes behind automation and the people issues associated with implementing it. He sees organisational maturity as a prerequisite for implementing effective (measured) process automation and ITIL as a useful framework for automated service delivery. He also looks after Collaboration and Business Process Management for Bloor, and takes a lively interest in the reinvention of the Mainframe as an Enterprise Server.

David has an honours degree in Chemistry, a graduate qualification in Computing, and is a Chartered IT Professional. He has a somewhat rusty NetWare 5 CNE certification and is a Member of the British Computer Society (he is on the committee of its Configuration Management Specialist Group).

He has worked in database administration (DBA) and operations research for the Australian Public Service in Canberra. David then worked for Bank of America and Swiss Bank Corporation in the UK, holding positions in DBA, systems development method and standards, internal control, network management, technology risk and even PC support. He was instrumental in introducing a formal systems development process for the Bank of America Global Banking product in Croydon.

In 1992 he started a new career as a professional writer and analyst. He is a past co-editor/co-owner) of Application Development Advisor and was associate editor for the launch of Register Developer. He helped organise the first London CMMI Made Practical conference in 2005 and has written for most of the major computer industry publications.

He runs his own company, **David Rhys Enterprises Ltd**, from his home in Chippenham, where he also indulges a keen interest in photography (he holds a Royal Photographic Society ARPS distinction).

Bloor overview

Technology is enabling rapid business evolution. The opportunities are immense but if you do not adapt then you will not survive. So in the age of *Mutable* business Evolution is Essential to your success.

We'll show you the future and help you deliver it.

Bloor brings fresh technological thinking to help you navigate complex business situations, converting challenges into new opportunities for real growth, profitability and impact.

We provide actionable strategic insight through our innovative independent technology research, advisory and consulting services. We assist companies throughout their transformation journeys to stay relevant, bringing fresh thinking to complex business situations and turning challenges into new opportunities for real growth and profitability.

For over 25 years, Bloor has assisted companies to intelligently evolve: by embracing technology to adjust their strategies and achieve the best possible outcomes. At Bloor, we will help you challenge assumptions to consistently improve and succeed.

Copyright and disclaimer

This document is copyright © 2018 Bloor. No part of this publication may be reproduced by any method whatsoever without the prior consent of Bloor Research.

Due to the nature of this material, numerous hardware and software products have been mentioned by name. In the majority, if not all, of the cases, these product names are claimed as trademarks by the companies that manufacture the products. It is not Bloor Research's intent to claim these names or trademarks as our own. Likewise, company logos, graphics or screen shots have been reproduced with the consent of the owner and are subject to that owner's copyright.

Whilst every care has been taken in the preparation of this document to ensure that the information is correct, the publishers cannot accept responsibility for any errors or omissions.

