



SIMPLIFYING INCIDENT AND PROBLEM MANAGEMENT WITH VBLOCK™ INFRASTRUCTURE PLATFORMS

IaaS COE
February 2011


WHITE PAPER

© 2011 VCE Company, LLC. All rights reserved.

Table of Contents

Executive Summary	3
The Challenge	3
The Solution	3
Problem and Incident Challenges in Traditional IT	6
Changing the Game with Vblock Infrastructure Platforms	7
Stateless Architecture	7
Service-oriented Computing	8
Management Tools	8
Consolidated and Correlated Event Views	9
How Vblock Infrastructure Platforms Address Incident and Problem Management Challenges	11
Visibility	11
Operational Silos	12
Incident Priority	13
Timely Resolution of Incidents	14
Flexibility	15
Conclusion	17
Next Steps	17

Executive Summary



When a system or process fails, the operational efficiency, productivity, security, and sometimes even the revenue of the business are impacted. When these incidents occur, the number one priority of the Information Technology (IT) Operations team is to restore business services to normal operations as quickly as possible. An effective incident and problem management program records, classifies, tracks, and closes all incidents in a controlled and consistent manner. How well an IT organization manages incidents to a successful conclusion is one metric by which IT's overall performance and value is measured.

In traditional data centers, the task of incident and problem management is often hampered by disparate architectures and point management tools that have limited visibility into the cause of the incident, what is affected, and what needs repair. Lack of knowledge about the resources that make up the affected service, their relationships, dependencies, and current state make efficient and effective management and resolution of an incident difficult, time-consuming, and costly.

Proactive incident and problem management programs combine services, technologies, processes, and data to minimize the impact of incidents on the business. Vblock™ Infrastructure Platforms provide opportunities for IT organizations to improve operational efficiency, reduce downtime, and increase customer satisfaction through more proactive incident and problem identification and management. The converged infrastructure, stateless architecture, and robust management tools enable IT organizations to fully visualize incidents and centralize management of them. By leveraging and integrating the incident and problem management capabilities of Vblock Infrastructure Platforms, customers can realize considerable gains in both hard operating expense efficiency and soft business revenue and productivity.

The Challenge

In a traditional data center environment, IT Operations are separated into silos of capabilities such as network, server, and storage. Although this traditional operating model allows for specialization and focus within these operational silos, it also propagates inefficiencies in many areas, including the effective identification, management, and resolution of business service incidents.

IT organizations face numerous challenges that impact efficient incident and problem management. These challenges include:

- Visibility
- Operational silos
- Incident prioritization
- Timely resolution of incidents
- Flexibility

The Solution

Vblock Infrastructure Platforms offer features and management tools that help IT gain valuable insight into the current state, resource relationships, redundancy, and dependencies of deployed business services. The Vblock platform features that facilitate effective incident and problem management are the stateless architecture, service-oriented computing, powerful management tools, and consolidated and correlated event views.

These features provide:

- Visibility to deployed business services across the entire Vblock converged infrastructure
- Flexible options for service continuation
- Alert and event coordination and correlation
- Reduced Mean Time to Identify (MTTI) the problem and Mean Time to Restore Service (MTTRS)
- Intelligent Root Cause Analysis (RCA) to help with service restoration and reduced redundancy events

Leveraging Vblock Infrastructure Platforms and their incident and problem management solutions help lower operating expenses, reduce outages, improve service levels, and increase customer satisfaction.

Scope

This paper examines how Vblock Infrastructure Platforms address the challenges faced by IT organizations in traditional data centers when identifying and resolving incidents and problems affecting delivery of business services. It examines the Vblock platform features that facilitate effective incident and problem management and the benefits provided by these features.

Audience

CIOs, Infrastructure and Operations VPs, Directors, and Managers, and any technical management responsible for defining IT strategy will benefit from reading this paper. Sales engineers, field consultants, professional services, IT managers, infrastructure architects, partner engineers, and customers may also find the content useful when integrating next-generation data center tools into their service management domain.

Terminology

The following table contains terms used throughout this paper.

Term	Description
Cloud computing	Cloud computing ('cloud') is where shared servers provide resources, software, and data to computers and other devices on demand, as with the electricity grid. Cloud computing is a natural evolution of the widespread adoption of virtualization, service-oriented architecture, and utility computing.
Converged Infrastructure	A converged infrastructure packages multiple IT components into a single, optimized computing solution. Components of a converged infrastructure solution include servers, data storage devices, virtualization, networking equipment and software for IT infrastructure management, automation, and orchestration.
Event	Alerts or notifications created by an IT service, configuration item, or monitoring tool.
Incident	Any event that is not part of the standard operation of a service and that causes, or may cause, an interruption to or a reduction in, the quality of that service.
Known error	A problem becomes a known error once a root cause for a problem is determined and development of a workaround is initiated.
Problem	A condition often identified as a result of multiple incidents that exhibit common symptoms. Problems can also be identified from a single significant incident, indicative of a single error, for which the cause is



Term	Description
	unknown but for which the impact is significant.
IT as a Service	IT as a Service (ITaaS) refers to a method by which an IT organization or service provider delivers well-defined services delineated at a clear business interface, usually to multiple consumers and/or different business units. Delivering services in this manner, as opposed to operating as pure infrastructure administrators, allows an IT organization to increase underlying resource utilization and reduce associated operational costs. It also enables IT to offer pay-per-use business consumption.

Problem and Incident Challenges in Traditional IT

The traditional approach to incident and problem management is essentially a sequential, labor-intensive, and time-consuming process of elimination of possible root causes. The process relies on the involvement and coordination of many human resources, some of whom may have other pressing issues in their queue. This coordination and interlacing of work items greatly lengthens the time to achieve a resolution. This in turn negatively impacts MTTRS, which directly affects service level agreements (SLAs) and operational costs to the business.

This section details the challenges to effective incident and problem management that IT organizations typically face.




Figure 1. Typical Challenges Faced by Traditional Data Centers

- **Visibility**—Operation of a traditional data center requires a great depth and breadth of IT solutions, which often results in disparate tools, overlapping information, and multiple layers of virtualization. With poor or incomplete visibility of critical business services at every level of the data center infrastructure, IT Operations teams struggle to reduce operating expenditures and consistently achieve SLAs.
- **Operational silos**—Technologies required to run the business often force IT personnel to specialize and excel in one area of the data center, such as database administration, network security administration, Windows client/server specialties, and so forth. Root cause identification and resolution usually requires IT administrators and experts from the various silos to spend hours and perform numerous checks against the various pieces of the infrastructure. Every hour spent on incident identification and triage increases the operational costs to IT.
- **Incident prioritization**—With IT budgets allocating fewer and fewer dollars for new projects, IT management is more concerned than ever with reducing operating expenses. The ability to effectively prioritize incidents is a critical component of efficient IT resource utilization. Without complete visibility of the affected services and the impact they have on business operations across the entire infrastructure, it is difficult to determine which incidents are more critical than others.
- **Timely resolution of incidents**—The IT support goal is to get from initial event to service restoration as fast as possible. The longer it takes to filter, analyze, and correlate alerts and symptoms to find the root cause of an incident, the longer the time to resolution. The ability to quickly correlate events and alerts is critical to effective incident and problem management.
- **Flexibility**—Traditional physical data center architectures often have one or more servers for a single service. When an outage occurs, the service is down until the server, database, or storage problem is resolved. Providing high availability through redundancy is typically the only option, but this requires

dedicating hot standby or clustered servers, which are expensive, operationally inefficient (servers just waiting to be used), and require additional management.

Changing the Game with Vblock Infrastructure Platforms

With Vblock Infrastructure Platforms, VCE delivers the industry's first completely integrated IT offering that combines best-in-class networking, computing, storage, virtualization, management, and security technologies with end-to-end vendor accountability.

Vblock Infrastructure Platforms provide many features that facilitate effective incident and problem management. These features are:

- Stateless architecture
- Service-oriented computing
- Powerful component management tools
- Consolidated and correlated event views through EMC® Ionix™ Unified Infrastructure Manager (UIM) Operations Center

This section examines the technologies and features that the Vblock platform offers IT organizations to meet the challenges of incident and problem management in the traditional data center.

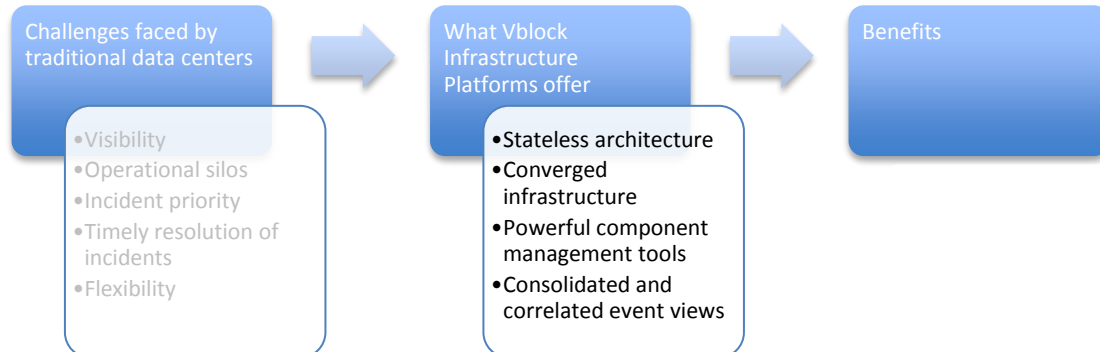


Figure 2. What Vblock Infrastructure Platforms Offer

Stateless Architecture

One of the most powerful features of Vblock Infrastructure Platforms is the stateless architecture. *Stateless* means there is no physical relationship between services deployed on the infrastructure and the underlying compute blades. Stateless architecture ties a business service to a service profile rather than any hardware. The service profile uniquely defines the service. Moving the service profile moves the service itself. Any service profile can be easily moved from one blade to another using UIM Provisioning Center. This flexibility allows IT to use a much wider array of possible workarounds to quickly resolve incidents and restore service.

For example, consider a VMware ESXi server running as a service on one Vblock blade. Utilization of the server has risen such that memory and compute performance needs to be increased to maintain service levels. One option is to re-allocate one or more of the guest operating systems to another VMware ESXi system, but this may overburden that resource. Another solution is to move the entire VMware ESXi server to an available blade with newer, higher performance CPUs and more memory. In a matter of minutes, users can have their hypervisor, guest operating system, and applications moved to newer, updated hardware, significantly reducing the maintenance downtime associated with updating hardware. Figure 3 illustrates moving a service profile.

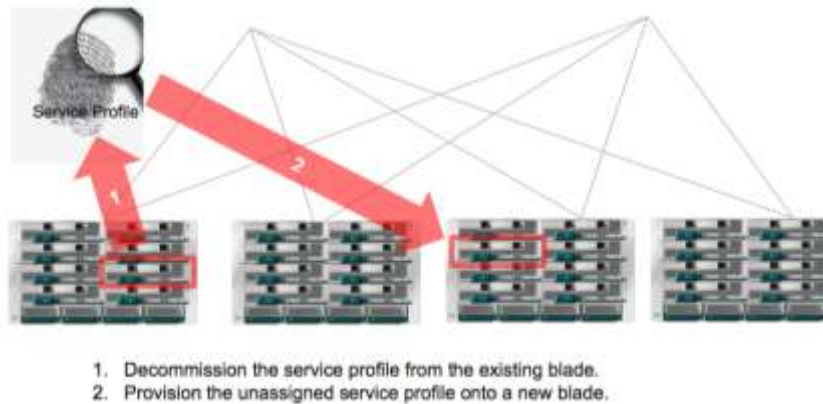


Figure 3. Moving a Service Profile

Service-oriented Computing

Vblock converged infrastructure is a natural fit for IT organizations that want to move to a more service-oriented management paradigm. In traditional IT environments, services are usually assembled sequentially piece-by-piece and then integrated into the existing management structure. The longer this assembly line method takes, the more difficult it is to see the service as an end-to-end solution. However, with Vblock Infrastructure Platforms, which provide a completely integrated IT offering for cloud computing, provisioning or changing a service takes minutes or a few hours instead of days and weeks. This consolidation, or abstraction, of multiple tasks into a single service request inherent in the Vblock platform allows IT organizations to more easily implement self-service portals for delivery of service-oriented compute services.

Management Tools

Because Vblock Infrastructure Platforms are so tightly integrated, they provide a unique, converged management opportunity centered on the concept of a complete service for IT organizations. EMC Ionix UIM management tools, such as UIM Provisioning Center, work with VMware® vCloud™ Director to deploy compute resources from the blade through to the actual requested application or service. These enhanced management tools provide visibility not previously available across the stack and control all the relationships built and maintained within and on the Vblock infrastructure.

Since these tools control and monitor changes in relationships, real-time relationship and dependency data is always available and accurate. Thus, what has always been a serious challenge to support staff trying to restore a service is now available as part of the Vblock Infrastructure Platforms solution. This paradigm shift significantly reduces the forensics efforts of support staff, thus reducing MTTI efforts, improving SLA achievement, and ultimately reducing operating expenses.

The benefits that the management tools bring to customers are clearly good practices: simplifying operational management of highly virtualized environments; lowering operating expenditures; reducing the time to resolve service-affecting issues; and ensuring SLAs are met. How these benefits are actually achieved in practice is critical to realizing IT staff efficiency.

The components of Vblock Infrastructure Platforms can be directly controlled using these independent toolsets:

- **Cisco® Unified Computing System™ (UCS) Manager**—Provides unified, centralized, embedded management of all software and hardware components of the Cisco UCS across multiple chassis and thousands of virtual machines (VMs). The entire UCS is managed as a single logical entity through an intuitive graphical user interface, a command-line interface (CLI), or an XML application programming interface (API). It provides flexible role- and policy-based management using service profiles and templates to reduce management and administration expenses, which are among the largest costs in most IT budgets.
- **EMC Unisphere™**—Web-based management interface for discovering, monitoring, and configuring EMC Unified storage. Unisphere offers quick access to real-time online support tools. It provides automatic event notification to proactively manage critical status changes
- **EMC Symmetrix® Management Console (SMC)**—Web-based management interface for discovering, monitoring, configuring, and controlling Symmetrix VMAX arrays. SMC enables initial system discovery and configuration, including device creation and configuration, along with basic device masking and support for managing local and remote replication activities.
- **VMware vCenter**—Provides unified management of all the hosts and VMs in the data center from a single console with an aggregate performance monitoring of clusters, hosts, and VMs. VMware vCenter Server gives administrators deep insight into the status and configuration of clusters, hosts, VMs, storage, the guest operating system, and other critical components of a virtual infrastructure.
- **UIM Provisioning Center**—Provides visibility across the service stack and information on the state of provisioned services.

These sources can be polled individually and/or can be integrated into a customer's existing management framework to provide significant intelligence about the resource relationships and dependencies, and the state of business services running on the Vblock converged infrastructure.

Consolidated and Correlated Event Views

The customer is still left with the formidable task of correlation and analysis of disparate events from multiple sources. The shift to highly virtualized, unified infrastructures naturally demands a similar shift in management capability. Consolidated incident and problem management for Vblock Infrastructure Platforms will be available with the release of UIM Operations Center (Figure 4).

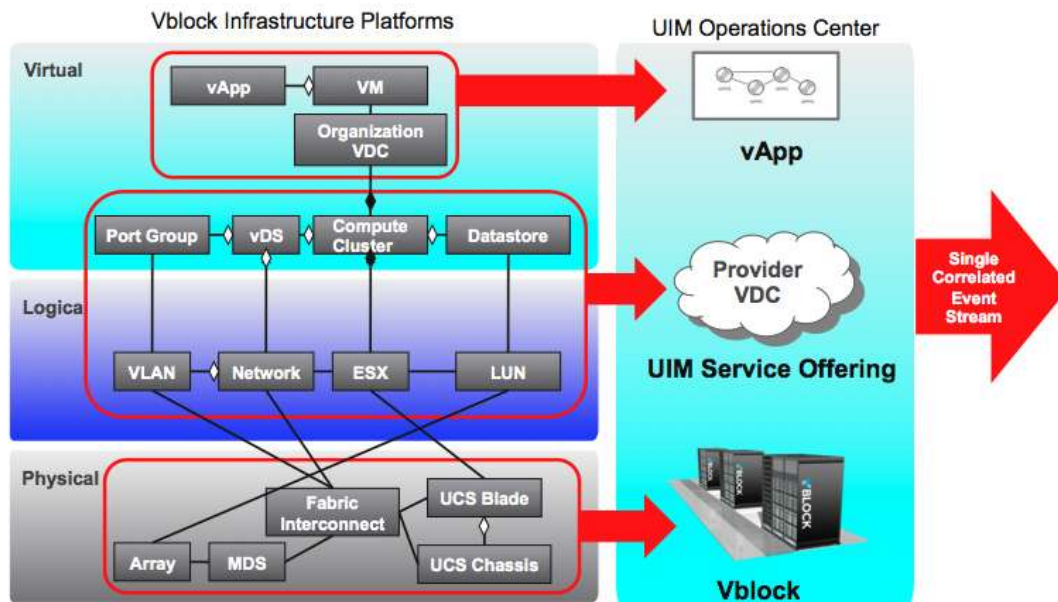


Figure 4. Unified Infrastructure Manager

UIM, which is included as part of the Vblock Advanced Management Platform (AMP), provides simplified management for Vblock Infrastructure Platforms by combining provisioning as well as configuration, change, and compliance management. UIM Operations Center pulls disparate data sources together to provide consolidated topology views, event correlation, and root cause and impact analysis. It is a single authoritative control point for operational management of the Vblock platform. UIM Operations Center delivers:

- Organized topology and events from
 - VMware vCloud Director
 - VMware vCenter Server
 - UIM
 - UCS Manager
 - Cisco MDS
 - EMC storage products
- Correlated information in context of higher order Vblock abstractions including
 - Service offerings
 - Provider and organizational multi-tenant virtual data centers
 - Virtual applications
- Root cause and impact analysis across the stack that provides critical information on availability, performance, capacity, and redundancy.
- Single console to organize, search, and display alerts and visual mashups.
- A correlated telemetry stream of topology and events to higher-order management systems through a RESTful web services API.

The capabilities of UIM provides customers with a correlated stream of events and an associated topology view that reduces operating expenditures, reduces time to restore service, and improves SLA metrics.

How Vblock Infrastructure Platforms Address Incident and Problem Management Challenges

Vblock Infrastructure Platforms represent a significant shift in operational management, as its management tools understand the relationships, states, and dependencies of converged architecture and offer IT Operations teams access to intelligence that can be easily integrated into existing service management frameworks. This section examines how the Vblock platform addresses the challenges faced by traditional IT organizations and provides resulting benefits from the move to a converged architecture platform.

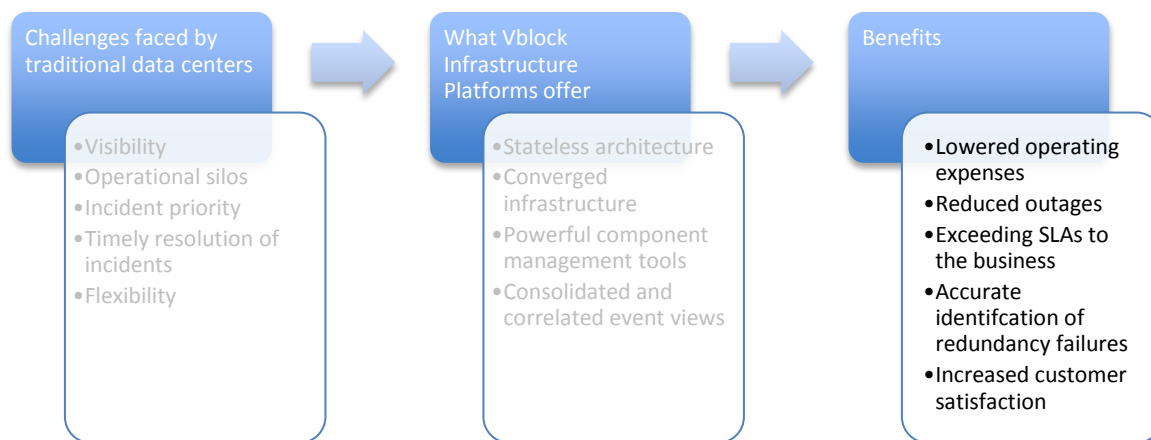



Figure 5. Benefits Achieved Through Vblock Infrastructure Platforms

Visibility

Visibility refers to IT's ability to see all the physical and virtual parts of the infrastructure that form a business service. When a problem arises, in order to identify what is happening in an affected service, IT needs to be able to construct the relationships between application and virtual machine, virtual machine and ESX server, ESX server and network, network and storage, and so forth. It can take longer to understand what is affected and what services are involved than it takes to actually fix the problem.

From a sheer numbers perspective, many IT organizations struggle to simply identify what assets are in their data centers and where they reside. Server sprawl, whether physical or virtual servers, makes inventory and mapping of applications to servers difficult. Silos of storage area network (SAN) and network attached storage (NAS) centralize storage management to some extent, but make it difficult to track which storage resources are tied to critical business applications.

Network security with edge firewalls and core virtual local area networks (VLANs) further adds to the complexity of understanding how and where data travels from client to server to storage and back. Without a solid foundation of tools to maintain an up-to-date inventory of assets, their mission in the data center, and their relationship to other assets (all of which are characteristics that can change at any time), reducing the time to restore service or the performance of a given service will remain a challenge. In other words, without visibility of the critical business services across the data center stack, operations teams will continue to fight upstream battles to reduce operating expenditures and consistently achieve SLAs.



How Do Vblock Infrastructure Platforms Address Visibility?

Vblock Infrastructure Platforms converged architecture provides a holistic view of the current state of running services. Management APIs provide up-to-date information on the status of Vblock Configuration Items (CIs) that can be populated and updated in a configuration management database.

EMC UIM Operations Center is the single gateway to Vblock events, by providing consolidated topology, event correlation, and root cause and impact analysis. It pulls together disparate data sources to provide a single authoritative control point for operational management of Vblock Infrastructure Platforms.

What is the Result?

Delivering pre-integrated management tools that have full visibility of Vblock Infrastructure Platforms, converged infrastructure can significantly reduce costs of operations. This reduces software costs and costs associated with developing custom integrations, automates event correlation and root cause analysis, and speeds the time to resolve issues, reducing the need for additional operations personnel.

Addressing visibility through the Vblock platform also results in more accurate identification of redundancy failures, improved SLA performance, and increased customer satisfaction.

Operational Silos

Operational silos challenge every IT manager. The depth of the technologies required to run the business often force IT personnel to specialize and excel in one specific vertical of the data center, such as database administration, storage administration, network security administration, application specialties, client/server specialties, VMware expertise, and so forth. With traditional IT infrastructure, there is little or no documented relationship between compute, network, storage, and ultimately, the application. When IT teams responsible for data center outages are not well integrated, operational silos lengthen the time it takes to resolve service-impacting outages.

MTTI, which is often identified as the longest stage of incident and problem management, is directly impacted by how long it takes to identify which organizational silo needs to be engaged to resolve the problem. Without some authoritative, correlated data source, such as a configuration management system (CMS) or database (CMDB) it is difficult to quickly determine which team to engage to resolve an incident.

Attempts have been made to solve this silo factor. However, these attempted solutions are often very expensive to implement and difficult to maintain, due to the challenges in understanding relationships across the stack in traditional IT environments – especially when those environments are changing daily from a technology and aging perspective.

How Do Vblock Infrastructure Platforms Address Operational Silos?

Vblock Infrastructure Platforms provide relationships that result in visibility of services across the operational stack. This helps IT quickly identify what components are involved with an impacted service, resulting in it taking IT Operations less time to verify if there is a problem with a component in their domain (such as a storage array, server, or network switch).

What is the Result?

Vblock management tools provide consolidated events about services that can help isolate which component is at fault. This ensures that tickets are routed to the correct team, which helps lower MTTI and MTTRS and reduces operational costs by not requiring all teams to come together to triage a problem. Addressing operational silos through the Vblock platform also results in improved SLA performance and increased customer satisfaction.

Incident Priority

With IT budgets allocating fewer and fewer dollars for new projects, IT directors and management are more concerned than ever with reducing operating expenses. The ability to effectively prioritize incidents is a critical component of efficient IT resource utilization. The last thing an operations manager wants to hear is that his most expensive administrator spent two days troubleshooting an outage that affected a non-essential internal web or blog site. Accurately prioritizing incidents from the outset is critical to focusing limited and expensive resources on the most important incidents and problems.

Determining which incidents personnel should focus on first can be difficult. If a server in an Exchange cluster is exhibiting performance problems, is that a Severity 1 or a Severity 3 issue? In this case, determining severity must take into account the role of the server; for example, is it a standby server, a database server for an internal blog site, or the core financial reporting database server for a business unit?

Even with very static services delivered on traditional, physical architectures, it can be difficult to establish the criticality of a performance or availability incident. For example, many customers have experienced alarms on servers where the staff could not easily determine what service was being delivered on the server.

How do you determine criticality? Without some way of understanding not only the component relationships, but also the current state of the infrastructure, it is difficult to avoid amassing all resources on any incident until criticality can be established. Unfortunately, by the time true business criticality is determined, expensive resources may have been expended on the incident, regardless of appropriateness. The need for understanding real-time state becomes even more imperative when you add virtualization with capabilities such as VMware vMotion, and with stateless architecture.

How Do Vblock Infrastructure Platforms Address Incident Priority?

Vblock management tools, such as UIM Operations Center, can provide IT Operations with the information it needs to know about the relationships between the business services being offered to their customers and the applications providing them, the servers (virtual or real) on which the services are running, and the infrastructures the services use.

IT Operations needs to know about the current production state of resources. Vblock management tools, such as Cisco UCS Manager and VMware vCenter, can provide information about the current state of the virtual and stateless components to higher-level management solutions like UIM Operations Center, which can then act as the single source of authority for relationships and state.

Redundancy and stateless profiling are central to the design of Vblock Infrastructure Platforms, as they make it easier to identify outages that are truly service affecting. The management tools built into the Vblock platform understand, track, and report on changes in the infrastructure, clearly understanding if an outage or degradation is related to redundant infrastructure or an available blade, or associated with a critical service profile. This frees Operations teams from having to develop this intelligence into their deployment, contributing to lower operational expense and an increase in overall visibility of the current environment.

What is the Result?

With an accurate understanding of relationships and the current state of Vblock resources, users can quickly prioritize the criticality of incidents that occur in the data center. This ensures that experienced personnel is focused on the most critical issues, lowers operational costs, and decreases MTTR.

Timely Resolution of Incidents

The timely resolution of incidents depends on quickly identifying a reduction in service, identifying the incident(s) (MTTI), determining an action path to resolution, and restoring service (MTTRS).

The time it takes to resolve the incidents experienced in a traditional IT data center is affected by more than operational silos and incident priority. Beyond those (admittedly significant) factors is the basic speed with which alerts and symptoms can be sifted, assessed, and connected to gain an understanding of where the true root cause lies. Here, once again, the difficulty in getting clear and direct visibility into today's complex data center environments makes this a slow and perhaps daunting task for IT staff.

Event Correlation

The goal of IT support is to get from initial event to service restoration as fast as possible through either a workaround or a fix. A long gap between event notification and resolution results in a longer outage and a greater impact on the business and SLAs.


A key factor in rapidly determining the significance of an event is the ability to quickly correlate events to the affected service(s). This activity is impacted by how effectively IT infrastructure and services are mapped and documented. The job of correlation becomes more difficult with less tightly integrated infrastructures, which is the case in most traditional IT environments. This is generally due to the widespread use of application-specific point monitoring tools that provide event data that is narrowly focused on only the infrastructure being monitored. Here are some examples of point monitoring tools and their visibility challenges:

Tool	Visibility Challenge
Server monitoring	Can report on CPU and memory utilization, but has no visibility into the storage or the virtualized environment running on it
Storage monitoring	Has no understanding of which applications, servers, or guest operating systems are using its capacity
Network monitoring	Can build network topology, but is challenged in understanding virtual application relationships

The challenge is to maintain an accurate view of the relationships and states of the infrastructures being monitored. Since point monitoring and management tools have limited knowledge of topology or state in the data center, IT Operations teams are forced to build or buy additional Manager of Manager (MoM) tools that attempt to correlate events from disparate monitors. This build or buy approach can be very costly and is extremely sensitive to changes in the environment; in addition, deploying and integrating these MoM tools often requires custom development through scripts and “glueware” – an expensive, error-prone, and time-consuming effort.

Root Cause Analysis

Once a series of events is correlated with a critical service incident, the IT support staff must leverage another set of tools to move the incident toward resolution as quickly as possible. When looking at MTTI, which is the most time consuming component of MTTRS, developing an accurate analysis of the root cause of the problem takes the bulk of the work.



Root Cause Analysis (RCA) provides IT Operations teams with the source of the problem. Many times there can be numerous downstream or ghost events that lead administrators or operations personnel on paths away from the real problem at hand. A true RCA process takes information on topology, relationships, and application dependencies and applies it to a rule base that can greatly reduce the time to find the real culprit of a service-affecting outage. This allows those responsible for service restoration to focus their energies on the true source of the problem and can significantly affect reducing MTTI.

RCA relies on detailed evidence in the form of events; it applies them to a knowledge database that maintains information on topology, relationships, and application dependencies. From this, RCA can quickly determine the source of the availability or performance issue. The accuracy of the knowledge database is critical for RCA to provide reliable results, which means that the sources of data for the knowledge database, whether CMDB or custom spreadsheets, need to be reliable, stateful, and reflect the accurate relationships and dependencies in the enterprise. Keep in mind that there are many challenges in pulling together accurate, detailed data on dependencies and relationships when deploying traditional, disparate architectures, because there is no single view of the environment.

How Do Vblock Infrastructure Platforms Address Timely Resolution of Incidents?

Vblock management tools provide a consolidated real-time view of business services, which gives the IT Operations team knowledge of the underlying converged architecture associated with a service and the information they need to make informed decisions. The ultimate goal is to empower level 1 support to resolve more issues themselves, speeding time to resolution and lowering operational costs. This ensures a much better chance of service restoration before escalation is required.

Vblock storage management tools provide more accurate reporting of events associated with critical business services, making it easier for support staff at every level to identify the real incidents associated with an affected service. This helps reduce MTTI, which is typically the longest variable in MTTRS.

The unique ability of Vblock Infrastructure Platforms to provide consolidated and correlated events, and to associate these with resource relationships and dependencies enables greater RCA efficiency. This significantly reduces the time to identify and restore from a service degradation.


What is the Result?

UIM Operations Center, which is a higher-level management tool, provides across-the-stack relationships of services and consolidation of events, facilitating RCA to greatly lower the time to determine a fix or workaround for a problem. This can significantly reduce MTTRS, as it provides the information needed to identify not only what has affected a service, but point to the “where” and “what” needed to resolve the issue.

The consolidation of events, correlation, and root cause analysis significantly reduces the necessity of the IT Operations team to develop and integrate these functions. The value of this unified infrastructure is that single point of operations intelligence that can reduce the time to restore an affected service.

Flexibility

Traditional data center environments add flexibility through fault tolerance, meaning the systems provide redundancy or high availability. Fault tolerant systems are a custom solution designed from the ground up that must be maintained out of cycle from other standard processes. They are costly, as they require duplication of servers and support for each business service. Organizations often do not have a tested process for providing fault tolerance; if they do, they are typically reserved for the most critical business applications. These extra resources cannot be shared with other services. The end result is that adding flexibility to traditional data center environments is an expensive, custom endeavor.



Traditional physical data center architectures often have one or more servers for a single service. When a service affecting outage occurs, the service is down until the server, database, or storage problem is resolved. Providing high availability options (redundancy) is typically the only option but requires dedicating hot standby or clustered servers, which are expensive and wasteful (servers just waiting to be used), and require additional management. Because of the overhead and cost of these solutions in both implementation and management, many customers reserve them for only the most critical of services. This has the following results:

- It limits options to IT Operations teams to leverage temporary workarounds, which could be used to restore most services (not just Tier 1) until a permanent fix can be implemented.
- It introduces challenges when tracking cluster or redundant hardware changes, as these solutions often require a unique or custom understanding of relationships and state of the infrastructure involved. For example, if an Oracle database administrator moves a database service from a primary server to a standby server during a maintenance window, how will IT Operations know that the service is up and there is just a reduction in redundancy taking place?

Many IT SLAs incorporate maintenance windows and unplanned service interruptions into their contracts, meaning that even though a service may be down for maintenance it is within the boundaries of the SLA. This affects the end user, who, regardless of the SLA agreement, cannot complete their jobs until the business service is available.

The optimal approach is to provide flexible fault tolerance across the data center, delivering a combination of redundancy and high availability to all services that need it. This enables organizations to keep services up during planned or unplanned downtime, allowing IT to negotiate SLAs that are more competitive and in the end, deliver a higher level of service to the end user.

How Do Vblock Infrastructure Platforms Address Flexibility?


Vblock Infrastructure Platforms provide three layers of flexibility to support redundancy and high availability:

- **Virtualization** provides features such as VMware vMotion and VMware vMotion Storage that provide high availability for virtual machines running on the Vblock platform
- **Fully redundant paths for network and storage** provide fault tolerance for all services running on the Vblock platform
- **Stateless profiles within Cisco UCS that keep blades independent from the ESX environments on them** allow for movement of a fully-deployed production ESX system to other blades, providing an economical N+1 redundancy capability. This is accomplished with minimal downtime and is available for any service running on the Vblock platform for planned or unplanned maintenance or hardware failures.

What is the Result?

These three layers provide organizations the ability to deliver higher uptimes, keeping end users in business and satisfied with IT's performance.

Conclusion



Vblock Infrastructure Platforms offer features and management tools that help IT gain valuable insight into the current state, resource relationships, redundancy, and dependencies of deployed business services, facilitating effective incident and problem management. Bringing network, compute, storage, virtualization, and management into one stateless package delivered with integrated, operational management tools allows for additional significant benefits:

- Lowered operating expenses
- Reduced outages
- Exceeding SLAs to the business
- Accurate identification of redundancy failures
- Increased customer satisfaction

The effort to integrate Vblock Infrastructure Platforms into the existing Service Management paradigm is reduced to leveraging the existing APIs and management tools that come with the Vblock platform, resulting in lowered operating expenses. Vblock Infrastructure Platforms purpose-built management tools, such as UIM Operations Center, are built with an understanding of the overall topology and relationships between network, compute, and storage. This provides faster identification of incidents, driving quicker resolution through correlation, root cause, and impact analysis, which allows IT organizations to consistently exceed service level agreements negotiated with the customer.

Next Steps

To learn more about this and other solutions, contact a VCE partner or representative or visit www.vce.com.



ABOUT VCE

VCE, the Virtual Computing Environment Company formed by Cisco and EMC with investments from VMware and Intel, accelerates the adoption of converged infrastructure and cloud-based computing models that dramatically reduce the cost of IT while improving time to market for our customers. VCE, through the Vblock platform, delivers the industry's first completely integrated IT offering with end-to-end vendor accountability. VCE's prepackaged solutions are available through an extensive partner network, and cover horizontal applications, vertical industry offerings, and application development environments, allowing customers to focus on business innovation instead of integrating, validating and managing IT infrastructure.

For more information, go to www.vce.com.



Copyright © 2011 VCE Company, LLC. All rights reserved. Vblock and the VCE logo are registered trademarks or trademarks of VCE Company, LLC. and/or its affiliates in the United States or other countries. All other trademarks used herein are the property of their respective owners.

