UPLOGIX

Meeting the Network Requirements of Cloud and Virtual Computing with Local Management

SYNOPSIS | The widespread adoption of virtualized computing (including as the enabling technology in "Clouds") and the various ways that enterprises will change to fully exploit it will have a profound effect on the network and on IT Network Operations Groups.

Customer expectations are changing. New requirements are emerging, service levels are becoming more stringent and some time-tested strategies for managing costs and ensuring adequate service levels are being invalidated. **Trying to use the same old network and network management strategies and tools will cause virtualization and cloud initiatives to fail or to incur runaway costs.**

Implementing Local Management will be necessary to meet customer expectations for virtual and cloud computing.

Climate Change in Virtualization

The idea of creating Virtual Servers, hosted on a Hypervisor installed on a single piece of hardware that manages access to physical resources and keeps them discrete, is neither new nor revolutionary in itself. It is, however, a very useful one.

> "What became truly revolutionary was the idea that virtualization's separation of the entire server from its underlying hardware could make it possible to treat pools of hardware as nothing more than an underlying foundation for a more self-managing and self optimizing system."

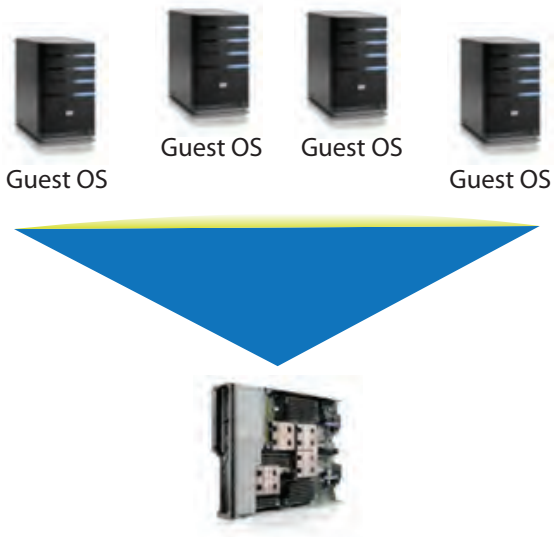

A simple Virtual Host with 4 Virtual Machines

The Virtualization of Servers increases platform choice by eliminating operating system hardware dependencies. As a result, enterprises can set up more easily managed homogenous hardware environments even when there is a requirement to host a variety applications developed for heterogeneous systems.

Virtualization also makes it much easier to provision and replace corrupted servers. As long as an existing server has available capacity, the IT department can quickly and easily create a virtual machine to meet a server requirement by an internal customer. This is a particularly visible part of what IT does for its internal customers and traditionally acquiring and configuring servers was a long and potentially cumbersome process, so the value of virtualization to IT operations in this respect was instant and obvious.

Finally, virtualization makes it much easier and safer to more fully utilize available hardware capacity. Companies can use VMs to control the portion of a physical server that an application will consume, allowing underutilized hardware to be ex-

ploited and also allowing for the purchase of more powerful servers, offering better price/performance for use by multiple applications.

These benefits make virtualization a very useful tool in the enterprise and have the potential to drive significant process changes in IT operations. However, what has become truly revolutionary is the idea that virtualization's separation of the entire server from its underlying hardware can make it possible to treat pools of hardware as nothing more than an underlying foundation for a more self-managing and self-optimizing system. When an application and its "server" are entirely software, then opportunities for more dynamic and constantly adapting automation of physical resource allocation to business servers and their associated business applications becomes possible.

In a fully virtualized datacenter decisions like which physical server to place a new virtual server on become fully automated, as does the process of provisioning servers to business users (self-provisioning). Virtual Machines can automatically relocate themselves according to specified rules if their current physical server is failing to meet performance requirements due to increasing application demands or technical issues with the hardware.

A Virtual Data Center

This changes everything, because it means that significant cost savings can be realized and economies of scale become a much greater factor in datacenter computing. Once a robust automation system is put in place, then as the datacenter grows the cost per unit of processing decreases much more than in the past, while at the same time reliability is increased.

Furthermore, to take advantage of this new cost-effective infrastructure, newer server-centric thin-client application architectures that move most application functionality to centralized datacenters are becoming predominant. This trend extends from traditional enterprise applications and even to things like centralized, served desktop operating environments.

Clouds Form on the Virtualization Front

A natural consequence of the introduction of technologies that provide economies of scale in datacenter operations is a rising interest in outsourcing. The idea of outsourcing IT functions has always been that an outsourcer could gain economies of scale by serving multiple clients, but these economies have traditionally been elusive.

Virtualization is not cloud computing and vice versa, but it is very unlikely that cloud computing and Cloud Service Providers (CSPs) would be generating the high level of interest in IT circles that they currently are if not for the ways that Virtualization enables them. Cloud Service Providers generally exploit virtualized data centers to offer their customers:

- **Elasticity** – the ability to provision and de-provision compute capacity very easily and usually automatically and instantly.
- **By-the-byte and subscription pricing** – the ability to charge only for actual usage. This also allows what would traditionally have been capital costs to be treated as an expense.

CSPs that only offer virtual datacenter services are often referred to as "Infrastructure as a Service" Providers, with a whole array of other services falling under the definition of "Cloud Service Provider," including companies that provide hosted development environments and hosted applications. Companies like Salesforce.com are considered a form of CSP.

Not surprisingly, many large internal IT organizations have embraced virtualization in order to provide their internal business unit customers with similar services and options, coining the term "Internal Cloud."

Network Operations Impacts

In many ways, the exploitation of virtualization and cloud computing services by enterprises threatens to take IT server and server management challenges and exchange them for network and network management challenges. At a minimum, the impact of the associated changes in technology managed, as well as the impact on business needs and expectations, will stress IT Network Operations groups severely.

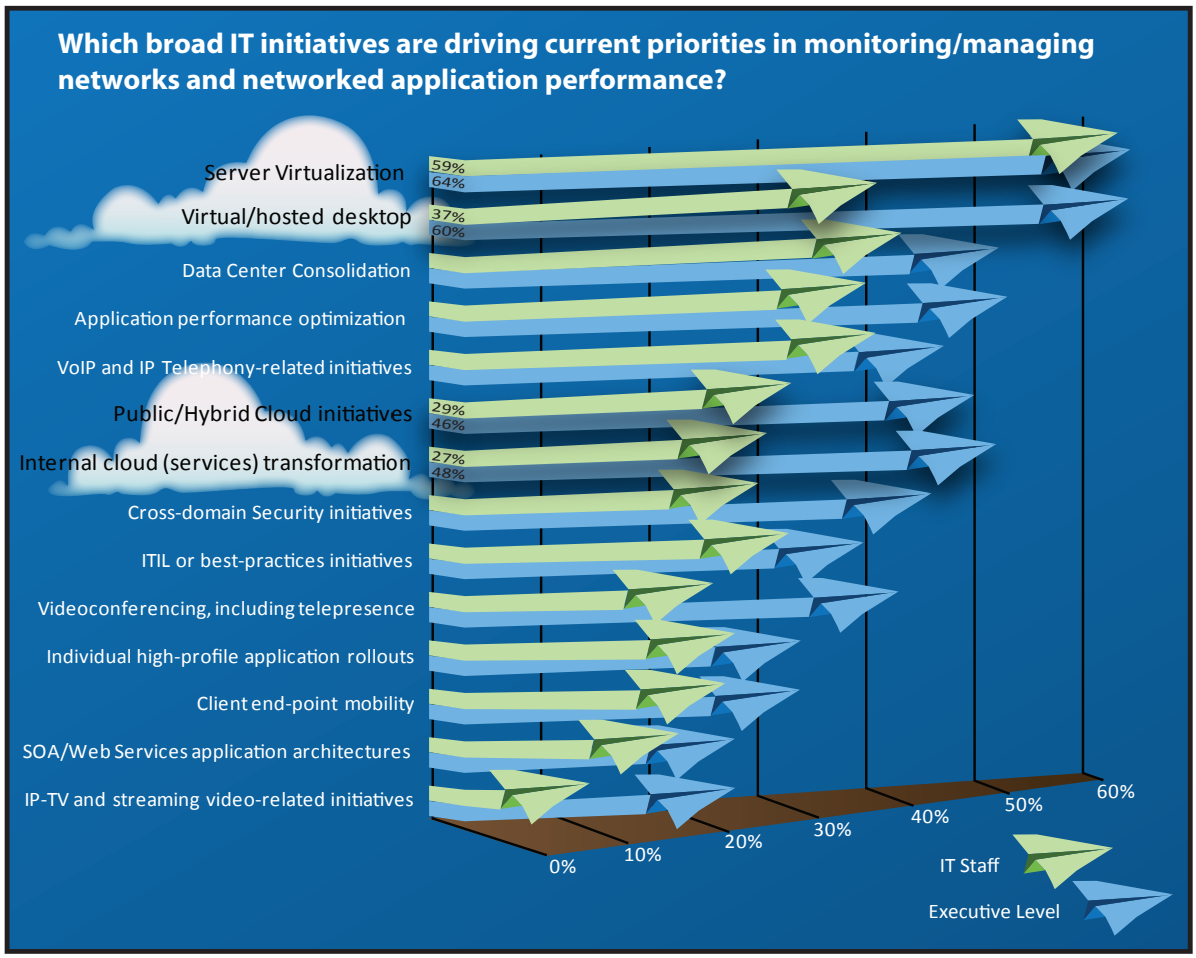

Research conducted by Enterprise Management Associates in 2011 shows that Virtualization is already the top driver of network management priorities.

We believe that there are four key ways that this will be the case:

- Cloud Service Providers are changing the ways that internal customers view IT services, raising their expectations and increasing their perceived options while increasing their reliance on the network.
- Highly centralized applications increase the utilization of – and the criticality of – branch networks and the WAN.
- In the virtualized datacenter, divorcing application architecture from physical architecture (a key step to achieving the sought after benefits) increases the criticality of every network device.
- The fact that every virtual host (hypervisor) has a fully configurable virtual switch inside of it dramatically increases the number of "network devices" that have to be managed.

Changing Business Expectations

Cloud Services Providers, typically by taking advantage of virtualization, have made it easier than ever for Line of Business users to rapidly and inexpensively source computing infrastructure, storage, hosted development platforms and enterprise class business applications over the Internet. Usually this can be done, or at least started, with a very low monthly subscription fee, no commitment and a credit card.

In most, but not all, enterprises, IT organizations are facilitating the effective exploitation of such services, establishing standards and processes and handling negotiations to make sure that costs and service levels are managed in line with business goals. At the same time, many IT organizations are gearing up to provide services to their internal customers on similar terms, also usually exploiting virtualization. This may include rapid or even self-provisioning of infrastructure resources and internal usage-based billing. In the vast majority of enterprises, IT services are comprised of a combination of both.

> "The necessary response is to establish local management nodes that complement and augment existing centralized management systems to address the most significant vulnerability of the new computing model – network dependence and sensitivity."

The most significant impact on IT Operations is the way that this trend is impacting internal customer expectations:

- **They Perceive Options** | Internal customers are coming to see internal IT as just one way to get the technology resources that they need. This makes it increasingly difficult for IT organizations to set service level objectives at a point that optimizes cost with outcomes based on business logic. Instead, CSPs, seeking to compete for revenue generating customers, drive the standard higher than might make sense.
- **They Act Like Customers** | Employees collaborate with IT to get what they need in the context of what is best for the enterprise overall. Customers

decide what they want and what they are willing to pay for it, then select a provider. Someone with a customer attitude tends to be concerned with getting the best deal for themselves and tends to demand more.

- **They Need the Network** | None of these services can be accessed without a robust network. Ironically, even as internal customers turn to third party service providers for a portion of the services that they once looked exclusively to internal IT for, their dependence on internal IT to enable them to access these services increases. Ultimately this shifts operational priorities toward the network. The more CSP services are used, the more sensitive to network outages and performance issues users become.

Centralization of Applications

In Client/Server application architectures, a significant portion of the processing typically takes place local to application users. This shields users to a large extent from intermittent WAN outages and congestion. In most client/server applications users can continue to work with varying degrees of diminished functionality indefinitely without access to the WAN, and in some cases with client heavy applications, even the LAN. Further, the quantity of data transmitted via the WAN is small as the data is typically limited to transmission of processed data such as database updates and as a result tends to be weighted heavily toward downloads.

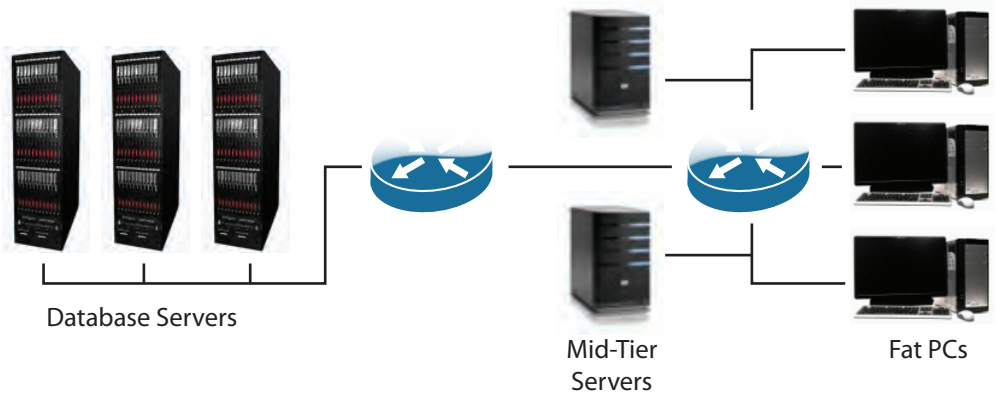

An Environment Built for Client Server Applications with PC Clients, Local Servers and a Centralized Database Back End

For an array of reasons, including lower management costs, ease of adoption and implementation and the availability of enterprise class third party offerings, enterprises are making increasing use of highly centralized applications. This includes web based SaaS offerings as well as internally developed or hosted applications

and application services. Current trends point to an extreme move in this direction, with Virtual Desktop Integration (VDI) becoming increasingly popular in large enterprises. VDI goes so far as to take traditional desktop OS functionality and centralize that in a virtualized data center along with application functionality.

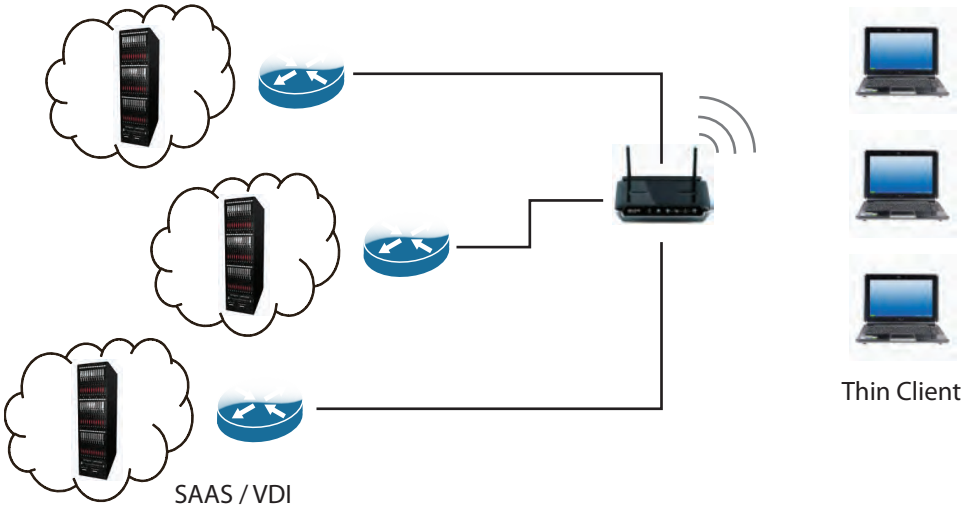

An Environment Built for Centralized Applications with Thin Clients
(e.g. SaaS Web Applications)

While there are a lot of good reasons for centralizing applications, this trend makes particular sense in light of the economies of scale offered by virtual data centers. Naturally, if infrastructure and management costs per unit of performance decrease as the data center scales, it make sense to move as much processing as possible into the data center. At the same time, this new strategy creates an array of new network management challenges:

- **Network Sensitivity** | User experience becomes extremely sensitive to network performance and availability. If the WAN or LAN goes down, so does the application, or in some cases with VDI, even the desktop, immediately and completely. If the WAN or LAN slows down, the application as experienced by the user slows down correspondingly and again, immediately.
- **Network Utilization** | In most cases network utilization will increase a great deal. Depending on how optimized the applications used are, it may be a little or it may be dramatic. In all cases the long-standing practice of optimizing network cost/performance by provisioning for more download bandwidth than upload bandwidth will rapidly become obsolete.
- **Monitoring** | Monitoring application and network performance from a user perspective to ensure service levels are met and even defining service levels becomes considerably more complex. In traditional application archi-

tecture, application performance could be inferred by monitoring database and middle tier server performance. In the new architecture, an application can be an array of application services hosted and delivered from multiple internal and external sources over different network segments that is never aggregated into the full application other than on the user's screen.

Physical Disassociation of Applications

In a traditional datacenter, application components from a single business application are installed on known servers connected by known network devices. From a network management perspective this provides two key advantages. First, it is possible to ensure that a specific set of properly configured network devices are used by the application. Second, the network devices used can be selected to be appropriate for the application they support in terms of performance and reliability (including HA) vs. cost.

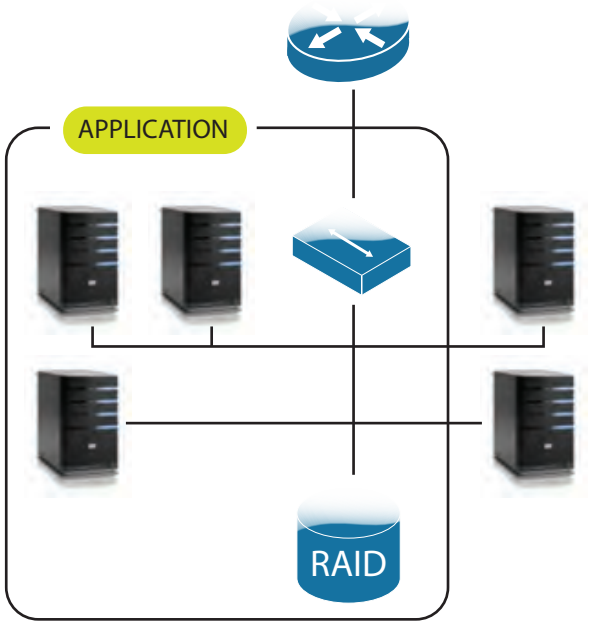

An Application Architecture Defined as a Specific Set of Software Components and Hardware Components

One of the primary benefits of virtualization is the fluidity of server resource allocation in response to business and application needs. Unfortunately, from a networking perspective this means that any business application component in a given cluster can be on any server in the cluster using any network devices used in that cluster and, in fact, any number of them. Further, VMotion, in order to deliver its benefits, requires that the system is constantly changing.

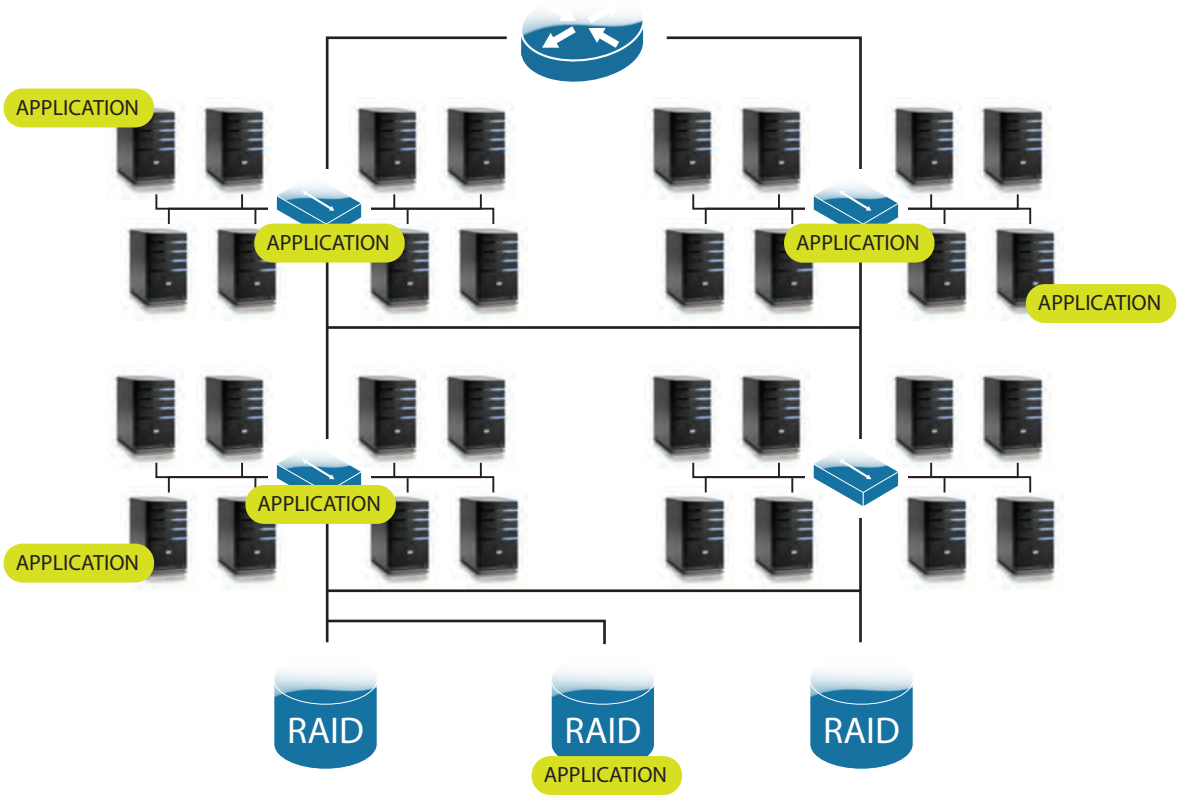

A Virtualized Application Architecture Defined as a Specific Set of Software Components with Virtual Server Requirements

As a result of this, every network device in the cluster must be treated as a critical component of the most critical application in the cluster. Furthermore, the potential impact of a single device failure can range from impacting a single (or even no) application to impacting every application in the cluster with a different impact at any given point in time. To some extent Mesh Network architectures can diminish this concern, but even where this strategy is deployed, there is no practical way to

ensure that single point device failures don't result in effective outages of critical applications due to resulting traffic overloads.

Virtual Switches Multiply Rapidly

Every single virtual host (the individual server with the hypervisor that the various virtual machines run on) requires a virtual switch to manage network communications between the VMs on it and between these VMs and the rest of the network. Each virtual switch requires almost the same configuration and management as a physical switch in the network and issues that impact physical switches can have the same impact on hosted VMs. Some of the challenges of managing an environment with many virtual switches include:

- **Sheer numbers** | In a virtualized data center the number of virtual switches to be managed can be an order of magnitude higher than the number of physical switches found in a comparable physical data center.
- **Ownership** | In many cases the default owner of the entire virtual machine, including its switch, is a system administrator. This creates an array of issues. System Administrators rarely have the expertise required to optimally configure switches or troubleshoot problems. Further, this exacerbates diagnosis when trouble occurs as discerning network performance issues arising from issues with the virtual switches from issues arising from the physical switches. These situations can be even more difficult than discerning server vs. network issues.

The Need for Local Management

Virtualization and cloud computing, along with the changes that are being implemented as a result, are driving much higher requirements for network performance and availability. They are also increasing the network operational management burden, including costs. Every network device is more critical, service level requirements are higher and there is more to deploy, configure, monitor and troubleshoot. Without changes in the network and how it is managed to anticipate all of this, the network will swiftly become the weak link in the chain and virtualization initiatives will fail as a result.

> "Local Management reduces the likelihood of problems occurring and emphasizes order of magnitude decreases in Mean-Time-To-Recovery when issues do occur."

Local Management

The only cost effective way to confidently meet the pervasive higher service levels that are mandated for networks by virtual and cloud computing is to implement a network management strategy based on Local Management that reduces the likelihood of problems occurring, and that emphasizes order of magnitude decreases in **Mean-Time-To-Recovery** when issues do occur. In this manner, Local Management provides a way to meet the elevated availability requirements of virtual and cloud computing that is cost effective enough to be a feasible element of any implementation strategy.

Local Management is a Network and Systems Management (NSM) practice that includes the following:

- A Local Management device or server co-located with and connected directly to the network gear to be managed (not using the network).
- Local Monitoring that takes advantage of direct connectivity to more rapidly identify failures and issues (typically the console port can support 30s intervals vs. 5-15 minutes with centralized network dependent monitoring systems using SNMP).
- Point-to-point network monitoring between Local Managers installed in proximity to both users and application service points of origination that can measure the performance of specific network segments. This makes it possible to accurately assess performance as experienced by users of centralized applications.
- Automation capable of performing a sequence of intelligent recovery actions in response to problems comparable to what would be done by an on-site technician.
- Configuration Management that automates complex configuration tasks to improve efficiency, reduce errors and allow for configuration changes to be included in automated problem resolution routines.
- Locally enforced role-based administrative access and detailed logging to prevent malicious outages or outages that result from non-compliance with standard administrative procedures and to protect against unauthorized data access.
- The ability to establish an out-of-band connection for when human intervention is required and to flow information to centralized network management tools when the network is down.

Organizational Impact

Under normal conditions network operations teams are stretched. More often than not, growing network infrastructures and the increased administration that comes

with them are expected to be handled with existing staff. IT organizations faced with the kinds of increased availability and service level requirements deriving from Virtual and Cloud computing will encounter a number of potential issues including:

- **Unacceptable Costs to Attain Acceptable Service Levels** | Given the magnitude of change in service level requirements, very few organizations will find it conceivable to staff up to the level that would be required to achieve them. This will be especially true for organizations with multiple remote locations.
- **Security Issues** | Given the duress that technicians will be operating under to respond to issues in a manner satisfactory to business users, the pressure to take shortcuts that circumvent network security policies designed to ensure data security will be much higher. Expediency will too often override security considerations at a time when data security breaches are becoming an increasing IT concern.
- **User Satisfaction Issues** | A transition to more centralized applications as part of a Virtual and Cloud computing strategy will introduce visible change to end users. If end users see increased outages and periodic slower performance that they attribute to these changes, organizational support for the initiative overall will suffer.

A Local Manager serves as a reliable, security policy compliant, network independent, on-site proxy for network administrators for a large portion of the tasks that they need to perform and reacts immediately when needed. Further, Local Managers can provide skilled administrators with appropriate permissions with a remotely accessible local toolkit for performing more complex tasks when greater expertise and attention is required. Local Management is a "force multiplier" for the network operations staff that will allow them to successfully overcome the Network Management challenges that Virtual and Cloud computing will present.

"Local Management is a "force multiplier" for the network operations staff that will allow them to successfully overcome the Network Management challenges that Virtual and Cloud computing will present."

High Availability Networking Cost & Management

One approach to meeting the required service levels could potentially be to implement a completely redundant network in order to increase the **Mean-Time-Between-Failures**. High Availability Networking can dramatically reduce outages and planned downtime as seen by end users and has been popular for mission critical portions of the datacenter for this reason. For enterprises moving to virtual and cloud computing that do not already have High Availability networks throughout their enterprise due to the especially high criticality of IT to their business (e.g. some sectors of finance, energy or the military), High Availability networking is not likely to be a practical approach to addressing the network requirements of the initiative.

First, as virtual and cloud computing becomes the norm in an enterprise, the network performance and availability requirements in the enterprise tend to increase everywhere, not just in select areas where mission critical applications are known to reside. Considering the pervasiveness of the requirement, the upfront capital costs would be prohibitive for most enterprises and would risk undermining the overall business case for cloud and virtual computing.

Second, the perpetual operational expenses associated with managing twice as many devices in a configuration requiring synchronization are typically 3X of a comparable non-redundant network. This strategy doubles the number of network devices to be provisioned, maintained, configured, upgraded, powered and tracked as assets and then deploys them in more complex configurations.

Considering this, High Availability networking is likely to remain as a partial solution for enterprises with a higher proportion of mission critical business processes. For most enterprises, considering the breadth of the requirement, full redundancy everywhere will not be a realistic option. In every case where redundancy is deemed necessary, there is also an excellent business case for the implementation of Local Management to reduce the costs of managing this larger and more complex network system.

Uplogix – Pioneering Local Management in the Enterprise

For years Uplogix Local Managers have been deployed on battlefields, at sea and wherever mission critical networks are maintained remotely. Today Uplogix serves many of the best known enterprises, delivering cost effective operations and the ability to meet stringent availability targets to organizations in energy, retail, finance, manufacturing, healthcare, government, education and other fields.

Uplogix is the first and only Local Management Platform. It is a patented solution designed from Day One to combine specialized software and hardware to provide a comprehensive Local Management solution.

To learn more about how Uplogix can help ensure that your network is more than ready to meet the demands of virtualization and cloud computing by cost effectively and dramatically improving the performance and availability of your entire network while enabling more efficient and less error-prone management of larger more complex environments, please visit www.uplogix.com/cloud.

ABOUT UPLOGIX // **Uplogix provides the industry's first local management solution. Our co-located management platform automates routine administration, maintenance and recovery tasks—securely and regardless of network availability. In comparison, traditional network and systems management depends on the network, uses multiple tools, and remains labor intensive. Uplogix puts the power of your most trusted IT administrator everywhere, all the time.**

Uplogix is privately held and headquartered in Austin, Texas with international offices in London and Monterrey. For more information, please visit www.uplogix.com.

www.uplogix.com | Headquarters: 7600B N. Capital of Texas Hwy. Suite 220, Austin, Texas 78731 | US Sales 877.857.7077, International Sales +44(0)207 193 2769