Efficiency and Cost Control In Network and Application Management

An ENTERPRISE MANAGEMENT ASSOCIATES® (EMA™) White Paper
Prepared for Fluke Networks

May 2009

IT MANAGEMENT RESEARCH,
INDUSTRY ANALYSIS AND CONSULTING

Table of Contents

Executive Summary

Networks are critical to business, and must be monitored and protected as would any other precious business resource. Today's connected businesses depend on reliable and predictable access to the application servers, data storage, and partner ecosystems which networks interconnect, and when there is any disruption, quick troubleshooting and rapid restoration is of paramount importance. The responsible approach to managing networks, however, dictates a shift from being efficiently reactive to problems towards a proactive, preventive disposition. Accomplishing this maturity step requires a comprehensive set of strategies, practices, and tools that reduces quality risk across the lifecycle of your networks and applications, focused on proactively assuring quality services in support of the broader organization.

The Proactive Management Challenge

Business today depends on the network. Applications are the lifeblood, and servers run the applications, but without the network that connects servers to the user community, those applications are useless as a business tool. We no longer ship tapes around for data exchange and we don't use hard-wired terminals any more. We depend on highly reliable networks for personnel and business partners to access the data and applications that are critical to the everyday flow of business operations.

> We depend on highly reliable networks for personnel and business partners to access the data and applications that are critical to the everyday flow of business operations.

When problems arise, they can be due to a wide range of issues. Assuring quality IT services means not only looking after availability, but also includes performance. IT end users expect to be able to not only reach their application, or get their VoIP phone call – it also has to be responsive and of sufficient quality to allow them to do their jobs. Outright failure issues are typically straightforward to find and fix – the more difficult and disruptive issues relate to performance. And in today's highly distributed, heavily virtualized, and multi-tiered service delivery architectures, there are many, many potential points of failure, either for breakage or as a source of performance degradations.

The risk, of course, is impact to the top line – brownouts or blackouts of access to critical IT services mean that business grinds down, or even stops completely, at significant cost to the organization. Downtime costs vary widely by size of organization and by industry vertical, and EMA research has seen it range from $25k to over $1 million per hour. What is much less well defined is the cost of degradation – the loss of efficiency and throughput when IT services are not performing up to snuff and all of the business processes that they support are subsequently slow. With this kind of cost on the line, it is imperative to focus not only on how to reduce the length of outages, but also on how to avoid them entirely.

Despite all best intentions, most IT operations teams today work in a predominantly reactive mode. Much of the time, they are made aware of problems only when notified of an incident by the service desk, or from end users directly. What commonly ensues is a fire drill, where personnel try to isolate, troubleshoot, and restore the service as quickly as possible. The business impact during this process depends on the amount of time that the service is impaired and the specific cost of outage of that particular service. While IT does not directly hold sway over the cost of outage of a service, they absolutely have control over the length of the outage. This time factor is commonly referred to as "Time to Restore" or "Time to Repair," and over time the average efficiency of this process is called MTTR, or Mean Time to Restore/Repair. Another important variable is MTBF, or Mean Time Between Failures, normally used to track time between outages, but which can also be used to track time between degradations, or any case where business processes are not operating at peak efficiency.

The challenge facing IT managers is how to reduce MTTR and extend MTBF. Doing either will protect the top line of the business. But what is often overlooked is that proactive approaches, those normally focused on extending MTBF, can also help the bottom line, by leveraging the value of preventative actions. As we all know, preventative measures, taken before services are actually impacted or are only marginally impaired, are almost without exception less expensive than reactive, firefighting actions which can tie up valuable human and technical resources across many IT groups, especially when dealing with subtle performance issues where root cause can be highly elusive.

> A more proactive approach towards network operations is fast become a mandate for IT executives.

As we all march steadily towards increasing reliance on network delivery of business-critical applications and services, the aspiration to achieve a more proactive approach towards network operations is fast becoming a mandate for IT executives. In the not-so-distant future, doing anything less could well be considered negligent to the mission of the organizations they serve.

Strategies for Proactive Efficiency and Cost Control

So how do you put yourself in the best position to assure quality IT services, in an efficient and cost-effective manner? Applications and services can suffer degradation because of many different issues – link congestion, client configuration, server health, storage delays, load balancing overloads, QoS policy mismatch, application design, unexpected routing or switching paths, and many, many more. If the source of a performance issue isn't immediately obvious, then the best place to turn as a next stop is quite often the network. Why start with the network? Because that's where all of the applications and services must travel, whether it is between tiers, between data centers, or between server and client. So the network is an excellent viewpoint from where you can localize problems and figure which domain-specific team should carry on deeper investigations.

In order to best make the transition from reactive to proactive operations, and from corrective actions to preventative actions as the *modus operandi*, you need to keep several key practice strategies in mind. First and foremost, it is essential to understand what is running across your networks, in terms of which applications and services are active and their normal patterns of use (a practice often referred to as "baselining"). Next, you need address the full lifecycle of your IT resources, including the network and the applications and services that the network will deliver. And finally, in order to de-risk all phases of your proactive management regime, you need to examine your management tools and technologies to make sure you have the best solutions you can afford. We'll look at each of these in greater detail in the following sections.

Expanding Visibility Up the Stack

Building proactive practices starts with understanding the big picture, and in this case means understanding exactly how the service delivery infrastructure is being utilized. It's not enough to understand aggregate traffic loads, variations, and peaks – you need more detailed and granular information on which applications and services are present, who is using them, what time of day they are normally active, how much traffic volume each one is creating, and the quality of user experience for each session. Only with this data will you be able to clearly prioritize actions when responding to an incident or taking preventative measures.

> Building proactive practices starts with understanding the big picture.

Key to this is recognizing all active applications and services, which means understanding a wide range of protocols, including those that are not immediately obvious, such as peer-to-peer. Other aggravating factors are compressed and encrypted traffic, such as secure http or virtualized desktop services – while the sources and destinations of this traffic can be readily recognized, the content cannot.

If you don't have visibility up the stack yet (meaning understanding not just network and transport, but application layers), the best place to start is with traffic flow records like NetFlow, J-Flow, or sFlow. These are generated by your network devices, and provide snapshots of essential traffic data including source and destination addresses, volume, and application in use for each session or transaction traveling across the network. There is little cost involved in turning this feature on; however, you will need some tools to collect, process, analyze, and present this information.

Flow records do have limitations, however, and often must be supplemented with other measurement/instrumentation types such as packet analysis and synthetic testing, in order to fully understand quality of experience (response times, in particular). These techniques also help with characterizing complex application types such as peer-to-peer, VoIP, Web2.0, and SOA. Further, these technologies can be the only means for revealing subtle application design flaws which adversely impact production performance.

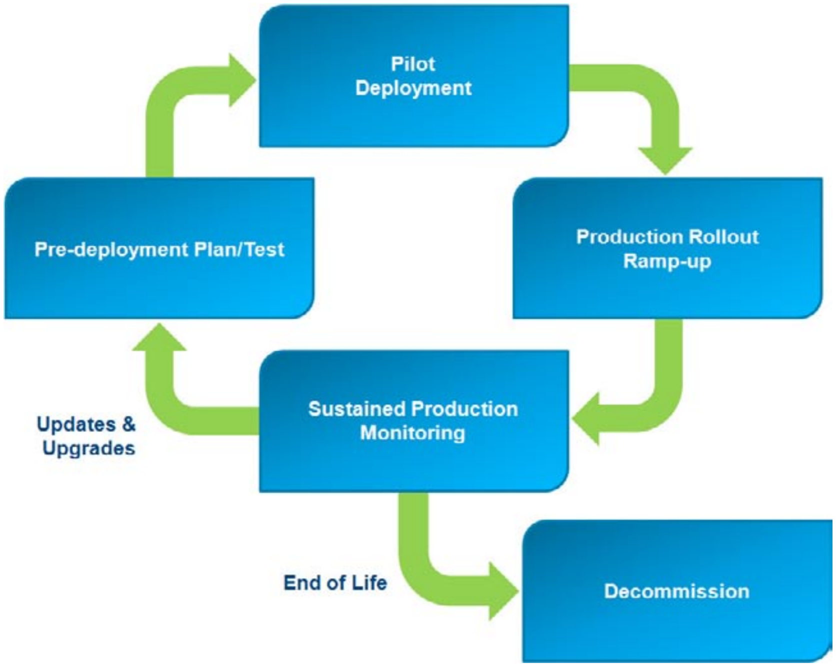

Managing from Cradle to Grave

In order to shift to a more effective, proactive set of practices, it is necessary to address the entire lifecycle of IT services and all of the contributing technologies that will be involved in delivery to end users. This means instituting standard practices for all phases from the very first planning through to eventual decommissioning. Following are some key considerations for each step along the way:

- *Pre-deployment:* The old saying is true – "an ounce of prevention is worth a pound of cure." What you do before a new application or service reaches the production network could save you many headaches down the road. If ever there was a time to be proactive, it is during this phase, where you can get ahead of the game. Use pre-deployment time to plan and prepare how you will assure quality in the live production environment, focusing on the following steps:
    - Define service objectives in terms of availability and performance, and decide what monitoring instrumentation points and metrics will be used to assess health.
    - Test new applications and services in a controlled lab environment, with an emphasis on characterizing behavior at nominal and high loads. Make sure to account for a range of latencies if the production environment will include WAN infrastructure.
    - Audit the existing production network to determine if it can support the new application or service, and plan and schedule configuration/policy changes accordingly

  Essential to ensuring successful practices at this phase will be working across political and organization boundaries, such as mandating that application/service developers include intrinsic measurability and network scale testing as part of the standard quality regime.

- *Pilot deployment:* In this phase, infrastructure policy and configuration changes should be put in place, and instrumentation should be trialed to monitor the pilot environment to assess availability and performance. Health reporting can also be trialed at this time and shared with the deployment team as well as upper management of IT and the affected business units.

  Though not ideal, it is not unusual for this to be the first time the network operations team may be seeing the new application or service, and if that is the case, then now is the time for characterizing behaviors and selecting monitoring regimes.

- *Production rollout:* Based on results of the pilot, adjust infrastructure policy and configuration profiles and monitoring plans for optimal results, and proceed to closely watch key performance indicators (KPIs) as scale is achieved. Service reporting should also be started on a regular basis during this phase.
- *Production monitoring:* At this point, if you've done the first three steps properly, you can turn the focus towards trending KPIs and watching for early warnings of situations that might impact availability or degrade performance, such as volume growth or slowing response times. If you recognize these signs proactively, before service quality is noticeably impacted, then you will have time to evaluate options for mitigation and make adjustments before the help desk starts getting calls. Often, this can include study of problem sources, with an emphasis on recognizing how the delivery infrastructure is being used, and looking for ways to wring more efficiency out of your existing infrastructure investments, such as usage policy changes or better job scheduling, rather than throwing more new hardware or bandwidth at the issue. If the preventative opportunity passes, because the onset of issues is very quick or an incident has simply eluded your proactive measures, then the focus turns to rapid troubleshooting and quick restoration.
- *Production updates:* Rolling out new versions, upgrades, and patches should be treated as a miniature version of a full lifecycle, with pre-deployment testing, piloting, and review/refresh of infrastructure configuration/policy, service metrics, and reporting practices.
- *Decommissioning:* An often overlooked phase, your management systems can help when it is time to move off of an old application or service, whether the task is discovering who is actively using the application or service or monitoring to ensure that vestiges do not continue to exist following turn-down.

The Right Tools for the Job

Hopefully the path to a new, efficient, proactive future is becoming clearer. What are the major hurdles on that path? And how do you shorten that path, and de-risk the transition that you'll need to make? The answer lies where you would expect – in the combination of people, process, and tools that you'll need to adopt. Fortunately, management technologies are progressing in scope, maturity, completeness, and integration, and can help to deal with the challenges you will face. Here are some key qualities and features of network management tools that you should consider when reviewing your current array as well as looking to improve your ability to become proactive going forward:

> Management technologies are progressing in scope, maturity, completeness, and integration, and can help to deal with the challenges you will face.

- *Intelligent, automated tools that apply (and capture) knowledge:* One of your best bets for improving efficiency and effectiveness within your operations group is to employ tools that have built-in knowledge bases regarding how networks and applications work. These typically take the form of "expert" or "advisory" features that can look at patterns of information and present operators with likely source scenarios, and recommended courses of corrective action. The two most common forms of these systems available today are packet analysis experts and root cause analysis engines. Packet analysis experts look at traces of network packets and apply rules and heuristics for revealing potential protocol mismatches or application design issues, among many other situations. Root cause engines help to localize the source of failures or degradations based upon a self-discovered understanding of the interconnected elements that make up your infrastructure.
- *Sustained monitoring for early warning:* Perhaps the most important step for making the move from reactive to proactive is to make sure that you are watching the service delivery environment on a constant, sustained basis. While sustained monitoring for fault and failure situations has long been the norm, less common is ongoing performance monitoring, and even less common is ongoing performance monitoring at the service and application layers. Some of the more innovative new technologies that have emerged over the past few years are intelligent analytics algorithms that can watch metrics collected via sustained performance monitoring and build an intelligent model of normal variations in KPIs over various time intervals. The models are then used to watch for unusual patterns and analyze them for relevance as early indicators of service quality issues, flagging those that are important for potential preventative actions.
- *Means and methods for easy collaboration:* Good data is only of limited use if it cannot be shared. Tools that facilitate collaboration can provide rich, configurable reports as well as means for easily distributing results across multiple groups within your organization. This is most often accomplished using some manner of Web-based reporting, and important to consider here is that the solution provide a method for ensuring appropriate levels of access, by role and/or domain, so only the information and control actions that are appropriate for each person engaged is allowed.

- *Ease of deployment and quick time-to-value:* While it has always been important to seek fast return on investment (ROI) for any expenditure, the pressure is greater than ever in the current economic climate. For management tools, this means products that can be deployed without substantial customization or services, and products that automatically discover and learn what is important in your managed environment. Also important are intuitive user interfaces and insightful reporting, so learning curves within both the operator and management teams can be as short as possible.
- *Consolidate solutions to cut the clutter:* There are a lot of management tools out there, and they all seem to be "best" at one thing or another. While it might seem an obvious choice to make sure you have the best tool for the job, in reality there are so many different functional capabilities needed to manage today's networks that a best-of-breed strategy can quickly become an administrative nightmare. In the extreme, EMA has witnessed an organization where over 150 different management tools were in simultaneous use. The licensing, maintenance, training, and interoperability (or lack of interoperability, specifically) in such an environment are major resource and efficiency drains. Your tool selections should cover multiple functions wherever possible – the trick is to make sure you prioritize your needs and be willing to accept solutions that might be second or third on your list for less critical capabilities.

A Real-World Example

Goal: A major credit union needed to monitor and troubleshoot its network in order to improve the reliability and speed of connections to its branch offices, and gain total network visibility from the data center.

Solution: Network monitoring and analysis tools deployed in the data center, with portable versions of the same monitoring system used for detailed diagnostics at potential problem origin points.

Results: A projected, cumulative five-year net benefit of over $800,000, realized by avoiding investing in new hardware, reduced maintenance costs, and improved productivity. The project had a payback period of three months. Troubleshooting and maintenance costs and time have been reduced, improving the bottom line. Most importantly, the bank is gaining more loyal customers because better network and application performance allows branch offices to provide improved services, bolstering the top line.

EMA Perspective

While the current economic hardships have sharpened everyone's focus on the bottom line, IT executives must be more mindful than ever of the need to identify strategies for improving their overall effectiveness in serving the business and organizations that depend upon them. This means looking for ways of getting more life and performance out of current infrastructure assets, improving staff efficiency, decreasing cost of operations, and supporting new revenue by accelerating deployment and removing barriers to success for new IT-enabled projects.

When the pressure to conserve capital and stretch personnel is on, management tools represent an important and constructive option. Many network and application management tools today offer rapid ROI due to their small cost relative to new capital infrastructure. They also represent an important hedge against operational risks arising from continually increasing complexity and the fact that businesses are becoming more (and not less) reliant on their network-connected infrastructures to conduct everyday operations.

Many network and application management tools today offer rapid ROI due to their small cost relative to new capital infrastructure.

Management tools can be instrumental in improving two specific categories of efficiency – responsiveness and problem avoidance. Rapid response can be realized via tools that provide automatic identification of the source of incidents as they arise and/or accelerated expert analysis and troubleshooting of those that are less immediately obvious. Problem avoidance can be accomplished by sustained monitoring of key performance and health metrics, and recognizing when changes in behavior indicate the need for preventative interventions. EMA advocates increased focus on this latter category as a best practice for enterprise IT. With the current state of management technologies available today, assuming a more proactive approach towards assuring IT services is more than simply possible – it's fast becoming a mandate for responsible operations.

About Fluke Networks

Fluke Networks provides innovative solutions for the installation and certification, testing, monitoring and analysis of copper, fiber and wireless networks used by enterprises and telecommunications carriers. The company's comprehensive line of Network SuperVision™ Solutions provides network installers, owners, and maintainers with superior vision, combining speed, accuracy and ease of use to ensure maximum network performance and the fast resolution of problems. Headquartered in Everett, Washington, the company distributes its products in more than 50 countries. More information can be found by visiting Fluke Networks' Web site at www.flukenetworks.com or by calling (800) 283-5853.

About Enterprise Management Associates, Inc.

Founded in 1996, Enterprise Management Associates (EMA) is a leading industry analyst firm that specializes in going "beyond the surface" to provide deep insight across the full spectrum of IT management technologies. EMA analysts leverage a unique combination of practical experience, insight into industry best practices, and in-depth knowledge of current and planned vendor solutions to help its clients achieve their goals. Learn more about EMA research, analysis, and consulting services for enterprise IT professionals and IT vendors at www.enterprisemanagement.com or follow EMA on Twitter.

Corporate Headquarters:
5777 Central Avenue, Suite 105
Boulder, CO 80301
Phone: +1 303.543.9500
Fax: +1 303.543.7687
www.enterprisemanagement.com