White Paper

Building an ROI: Application and Network Performance Management

Believing there are benefits to improved network and application performance is not enough for most organizations. The ability to quantify cost savings, improved productivity or reduced risks is a critical component in justifying an investment in application and network performance. This white paper will list eight key areas where cost savings can be quantified. Each organization will have different results and savings – some savings might be spread out evenly while others will be skewed to only one or two criteria.

FLUKE networks

Table of contents

Introduction

Enterprises of all types and sizes – across a wide array of vertical markets – are increasingly relying on the network and the applications that traverse it for organizational efficiency and business productivity.

Consider these examples:

- A global conglomerate rolls out an ERP systems connecting hundreds of locations running over its network.
- A Fortune 500 company determines what to manufacture at its plants spread across the globe based on daily updates to its sales forecasting applications. Ninety percent of the individuals entering these forecasts are in branch offices, whereas a decade ago all sales personnel worked out of corporate headquarters.
- An enterprise decides to roll out Voice over IP (VoIP) across an existing network for additional features and cost savings.

As these examples illustrate, enterprises are using new technologies and applications as a competitive differentiator as well as a tool to reduce costs and/or improve productivity of customers and end users. With the emphasis growing more on applications and networks, the risk of degradation or downtime is more critical than ever.

Enterprises today can "see" the value of improving network application integrity and efficiency via proactive monitoring and management of applications and the infrastructure. However, "proving" the value of improved performance and efficiency by quantifying savings or reduced costs through a positive return on investment is much more challenging for many enterprises.

There is a huge difference when it comes to qualitative versus quantitative assessments, as seen in the accompanying table, **Qualitative Assessment vs. Quantitative Assessment**. Typically, individuals within an enterprise focus on the qualitative point of view instead of quantifying the impacts of improved services and reduction in financial risks or exposure. This table highlights the differences between the two approaches using an improvement in network application performance and efficiency as an example.

Qualitative Assessment	Quantitative Assessment
The operations center will be more efficient in doing its day-to-day job with the new system.	The new system will reduce the number of trouble tickets by 45% and the time to resolve an issue by 55%, resulting in a total cost savings of $124,575 this year.
The network engineers will take less time to configure and fine tune our new MPLS-based network.	On average, this system reduces the configuration time (both initial and fine-tuning) by 70%. With an average of one hour per configuration and two hours for change management at each of our 80 sites, multiplied by the average $50 per hour salary of our network administrators, we will save $8,400 annually.
The new system will help us identify and resolve potential virus attacks faster which reduces cost.	Historically, a virus attack takes an average of 6 hours to identify and stop but with the new system, it will only take 2 hours. For our enterprise, the cost of a critical application being down is approximately $11,000 per hour, thereby saving $44,000 per virus attack. With an average of two attacks per year, the annual savings is $88,000.

It is virtually impossible to calculate a return on investment – either positive or negative – with qualitative assessments only. This paper will focus on how you can identify important areas of improvement through increased network application performance and efficiency. More importantly, you will learn how to quantify the potential savings gained by enhancing network and application performance management.

Improving network application integrity, performance and efficiency

An enterprise's reliance on its network and applications is likely to continue growing exponentially. Revenues, compliance, inventories, and customer support are inextricably linked with network and application performance. While the business-critical applications running over the network vary by enterprises – ranging from SAP and Oracle to Voice over IP (VoIP) and home-grown applications – the exposure of poor performance can mean bottom-line impact ranging from the thousands to the millions of dollars.

IT managers require the effective, flexible tools that can provide the highest degree of visibility across the infrastructure. Today's network managers are no longer responsible for only bandwidth and connectivity issues, but must take into account how the network impacts both applications and users.

Fluke Networks' Visual Performance Manager is an industry-leading solution for managing, troubleshooting and optimizing application and network performance. By providing an integrated view utilizing different data sources, Visual Performance Manager a complete view of application and network performance – so IT departments can take a holistic approach to network and application performance management. The objective of the system is to enable IT departments to measure the success in business terms based on the high performance of the mission-critical applications.

The objectives include:

- Increasing application availability by improving business performance with a detailed understanding of both the applications and network domains.
- Reducing operational costs by developing more efficient application performance management solutions and by not resorting to always having to add more bandwidth.
- Controling potential problems by proactively accessing performance data from across the entire infrastructure to uncover unauthorized applications, bandwidth hogs, viruses, congestion areas and trends.
- Optimizing bandwidth by allocating the right amount of bandwidth to maintain application performance without exceeding budget limits.

The remainder of this paper will demonstrate how to quantify cost savings and reduce exposure in each of these areas. This paper assumes you see the value improved application and network performance – but you need some way to quantify the savings so you can justify the budget for the solution. As you read the examples below, think how these issues have impacted or may impact your organization going forward.

Quantifiable impacts of better network efficiency

When quantifying the impact of improved application integrity and network efficiency, there are generally two types of savings – hard costs and soft costs. Hard-cost savings are tied directly to the amount of money saved from the impact. These can include reducing circuit size due to over-provisioning, lowering the number of trouble tickets, eliminating overlapping tools, and reducing the risk of lost revenue. Soft-cost savings are typically tied to improved staff productivity so employees are more efficient and can do more with their time. Both hard- and soft-cost savings must be accounted for in quantifying the impact of improved network efficiency.

As highlighted earlier, proactive monitoring and management of critical infrastructure and applications can greatly increase the success of an enterprise.

Areas of impact include:

- Increasing revenue and production capabilities
- Optimizing bandwidth expenditures
- Decreasing the number of overlapping tools
- Improving the identification and repair of virus and worm attacks
- Reducing application mean time to repair (MTTR)
- Lowering the number of trouble tickets and support calls
- Reducing network configuration time
- Improving end user productivity

This paper will drill down into each of these areas of impact in more detail as well as provide an example of how to quantify the impacts of improving network and application efficiency.

Increasing revenue and production capabilities

With the increased importance of the network and the applications traversing it, every enterprise must be able to answer the following question, "If my most business-critical application is down for a period of time, what is the cost per hour to my organization?" These costs can vary quite a bit depending on the enterprise. A financial institution might lose $500,000 in transactions for every hour of downtime during Wall Street trading hours. A wholesaler might incur penalties of $50,000 per day if a shipment is delayed to a vendor. A manufacturing company might lose $250,000 by having to shut down a production line because the just-in-time inventory system was taken offline.

Each of these scenarios has three components – the area of the business impacted, the financial risk, and the duration of the incident. Improving network application integrity and efficiency only reduces the amount of time the business is impacted, but the value is immense. Since the overall exposure to the enterprise is compounded by the time and financial risk, the reduction in the time impacted provides substantial savings. A system that either eliminates a percentage of negative impacts or reduces the MTTR provides a hard-cost savings for the enterprise.

When credit card usage stops

A national retail chain with 1,800 stores allows customers to use credit cards for purchases. On average, a store processes $23,250 in credit cards per hour. If there is an outage for the verification, the chain has two options: deny credit cards or accept credit card purchases with no approval. In today's world of decreased reliance on cash, denying the credit card is not a viable option for this enterprise. Historically 1.25% of all credit card transactions should be denied due to bad, stolen or over-credit limit criteria.

If this chain has a single site outage for 2.5 hours, there is exposure by having to resort to accepting credit cards without approval. There is a $726.56 impact in this per location ($23,250 x 2.5 hours x 1.25%). Now, $726.56 may not sound like a great deal of money, but when 10% of the locations could not approve orders, the exposure is now $130,781 ($726.56 x 180 locations).

With improved network efficiency and troubleshooting capabilities, the retail chain was able to sectionalize the problem and reduce the time to repair by an hour (by 40%). The hard cost savings to this retail chain would be $52,313 ($130,781 x 40% time savings).

Optimizing bandwidth expenditures

The cost of bandwidth can consume as much as two-thirds of total networking budgets for many enterprises. IT organizations must walk a fine line between having sufficient bandwidth resources for business-critical applications and not wasting critical budget dollars by over engineering the network infrastructure.

Compounding the issue for many enterprises is the proliferation of bandwidth-consuming applications across the enterprise. Enterprises are rolling out new applications that are impacting existing networks including VoIP, Oracle, SAP and Citrix. When enterprises deploy new applications, many times they add bandwidth not knowing if and where they actually need the increased resources. With tighter IT budgets, the days of throwing bandwidth at issues such as these are numbered.

Upgrades, only when necessary

A medium-sized organization is planning on rolling out VoIP across its 22 domestic locations. Due to the perceived utilization of the existing network and the delay-sensitive nature of VoIP, the enterprise decided to increase the speed of the 18 fractional T1 sites by 256Kbps and add a second T1 to the 4 other larger sites. At an average cost of $1,100 per month per T1 circuit and $340 per month for the increase of 256Kbps, the total annual cost of the upgrade would be $126,240 [$1,100 x 12 (for the number of months) x 4 (for the number of T1 sites), plus $340 per month x 12 (for the number of months) x 18 (for the number of fractional T1 sites)].

With a tool that provides detailed utilization, an enterprise can "right size" the circuits – having sufficient bandwidth for critical, delay-sensitive applications without wasting resources. In this scenario, of the 18 sites that were getting an upgrade, only six really needed 256K more bandwidth, four needed 128K (at a cost of $190 per site) and eight didn't need an upgrade. Of the four larger sites, one needed the T1 upgrade and the other three only needed a 256K upgrade. So instead of spending $126,240 annually on upgraded bandwidth, the enterprise only needs to spend $59,040 ($1,100 x 1, $340 x 9, $190 x 4 per month x 12 months). That is a cost savings of $67,200 or a reduction of 53.4 percent.

Eliminating redundant tools/equipment

An easy way to calculate savings that goes straight to the bottom line can be obtained by eliminating redundant tools or CPE including equipment leases and maintenance. If an enterprise can eliminate equipment such as portable protocol analyzers or "dumb" CSU/DSUs, the leasing and/or maintenance cost can be avoided. This equipment is no longer needed because the application and network performance management system provides similar functionality.

Reduce overlapping tools

A network application integrity tool allows an enterprise to eliminate 30 T-1 and fractional T-1 traditional CSU/DSU units across the infrastructure. The monthly leasing price for each device averages $80 (equipment and maintenance). The new solution that terminates the circuit allows the enterprise to eliminate the legacy DSUs saving $28,800 annually (30 DSUs at $80/month x 12 months).

In addition, the new tool provides extensive application visibility, so the enterprise is less reliant on portable protocol analyzers. Per device, annual maintenance costs approximately $4,500 per year (average price of $25,000 per device x 18% maintenance cost). By eliminating three portable analyzers, the savings is $13,500.

Improving identification and repair of virus and worm attacks

Even though every enterprise has firewalls, virtually every organization has been penetrated to some degree within the past year by a virus or worm attack. Enterprises tend to classify virus attack impacts in two ways: severe threat and residual impact. The severe threat is the most dangerous. Severe threats degrade the applications and network to such an extent that key programs or business activities are impacted, leading to possible lost revenue or incurred costs. The residual impact occurs once the severe threat is over but the enterprise must still update its systems and perform maintenance. On average, critical threats tend to last approximately 3-6 hours while residual impacts tend to last several days

The cost of a virus attack

Over the past year, this enterprise has experienced five virus/worm attacks with a varying degree of pain. The severe threat length ranged from two hours for a smaller, common security breach to nine hours for a severe attack that crippled the organization for over a day. For this enterprise, the cost of application and network downtime/degradation has been calculated to $9,000 per hour (this amount is impacted by criteria such as the size of the organization, the number of users and the importance of application and network integrity).

Determining a savings for improved virus/worm detection and resolution is based on how quickly the attacks can be solved compared to today's averages. To be conservative, this enterprise only wants to factor in the severe threat and not include residual impact. On average, the five attacks over the past year took two hours to eliminate the severe threat. With the network application performance solution, the amount of downtime/degradation should be reduced 70% by identifying the spike caused by the attack and using the troubleshooting tools to isolate and solve the problem. In this scenario, the savings is calculated to be $63,000 (1.4 hours saved x 5 events x $9,000 per event).

Reducing application MTTR

For most organizations, poor application performance is likely the most difficult and troubling problem. There are so many factors that can cause poor application performance, ranging from the local loop to the circuit to the network to the application servers. And if the problem is intermittent, the challenge grows exponentially.

The key to reducing application MTTR is the ability to isolate where – across the wide spectrum of possible causes – the problem is occurring. More important is the ability to go back in time, whether an hour, a few days or a couple of weeks, to easily identify what caused the intermittent problem so it does not return and pose greater threats later. The goal for enterprises should include identifying degradation in application performance before end users are impacted, thereby reducing the risk to the organization.

Solving the hard-to-isolate problem application

This calculation is similar to the virus/worm attack scenario above, but this one deals with typical day-to-day issues in identifying and resolving trouble tickets like "I can't process the order" or "I can't get to e-mail." For this enterprise, problem identification and resolution are harder because it isn't a single catastrophic issue causing the performance issue. For these types of intermittent issues, the IT organization focuses on when users cannot use the applications completely and estimates approximately six hours to resolve these issues. This does not include poor performance where applications may be slow but can still get through. The cost of poor application performance is $6,000/hour for this enterprise.

By having a robust troubleshooting tool which provides extensive Layer 1-7 visibility and both a real-time and a graphical "back-in-time" view, the amount of time saved in solving the problem has conservatively been estimated at 50%. For the savings calculation, the enterprise experiences this severe impact once every two months. The savings would be $108,000 annually for this scenario by reducing application MTTR (3 hours saved x $6,000/hour x 6 impacts).

Lowering number of trouble tickets and support calls

Many IT organizations are deluged with trouble tickets and support calls from end users. Handing the individual tickets and calls can be very time consuming and expensive. In addition, the reactive nature of troubleshooting tickets typically means other end users will be impacted by the same performance degradation. This creates a lose-lose situation where end users are negatively impacted and IT incurs additional costs in trying to resolve the problem.

Quantifying the savings of reducing the number of trouble tickets and support calls is one of the easiest calculations for determining savings. Enterprises traditionally know the number tickets/calls per month, the length of time per ticket, and the cost of support staff. Those are the critical figures needed to quantify the savings of improved network application performance – by lowering the total number of tickets opened and reducing the length of time to handle each ticket.

Reducing trouble tickets and MTTR

Over the past year, this enterprise has averaged 200 trouble ticket calls per month. Historically, a ticket takes 30 minutes to open/handle and the average IT support staff hourly wage is $30 per hour. By improving application and network performance and providing critical visibility to resolve trouble tickets faster, the enterprise assumes there will be 25% fewer tickets opened, and those tickets that are opened will be resolved in one-third less time than traditional standards. By reducing the number of tickets, the enterprise would save $22.50 per ticket not created ($45/hour IT loaded salary – $30/hour plus 50% for benefits – x 1/2 hour per call). Over the course of a year, the savings of lowering the total number of calls is $13,500 (50 calls/month x $22.50 savings x 12 months). In addition, the reduction in call duration could save even more money. The 1/3 reduction of 30-minute call would save 10 minutes per call. The savings would be $6,750 ($45/hour IT loaded salary x 1/6 hour savings x 150 tickets x 12 months). In this scenario, the trouble ticket savings is over $20,000 annually.

Reducing network configuration time

There are two major challenges in relation to network configuration – ongoing and network refresh. For ongoing network support, typical moves/adds/changes require configuration changes such as remapping virtual circuits or re-configuring class of service (CoS) settings. Minimizing the amount of time and resources required for ongoing configuration is key to reducing cost.

A network refresh typically places a larger burden on IT staff because each and every location will likely need to be configured initially and some percentage will need to be re-configured or fine-tuned over time. The deployment of MPLS-based services will likely drive configuration times higher when CoS settings are deployed. Historically, many enterprises require

at least three configuration attempts for initial set up and semi-annual fine-tunings of the CoS settings. Each configuration or fine-tuning pulls resources from the network organization.

A network application performance management solution can reduce the amount of time and the number of configurations needed for ongoing maintenance and network refreshes. By being a point of demarcation and completing active connectivity tests, the configuration maintenance is made easier and will reduce the amount of time required. In addition, MPLS-based networks with CoS settings will both greatly reduce the amount of time for set up but reduce the number or re-configurations required to optimize the applications and network.

Making network configuration easier

In this scenario, the enterprise averages one ongoing network configuration per site per year. A typical configuration requires 90 minutes of IT support staff. With the tools to test connectivity remotely, the time required can be reduced 80 percent. With 70 sites, this enterprise could save $6,120 by reducing the configuration time (70 sites x 72 minutes saved/site x $60 loaded salary – $40/hour + 50% benefits).

Quickly adding to the savings would be the re-configuration and fine-tuning of complex networks including MPLS-based networks with CoS. On average, enterprises require three configurations for CoS set up and at least two fine-tunings of CoS per year. Assuming the initial configuration is completed in a single step, the savings are $12,600 (2 less configurations x $60/hour loaded salary x 70 sites x 1 1/2 hour per site). In addition, with detailed reporting on CoS settings, the amount of time for fine-tuning is reduced by 85%, which saves $10,710 (70 sites x 2 fine-tunings x 1 1/2 hours per site x $60/hour loaded salary x 85% time reduction).

Improving network user productivity

A challenge for many enterprises is quantifying the impact of poor application and network performance for end users. While it is simple to say, for example, end users will be impacted if the order processing system is down for 6 hours, it is much harder to precisely quantify the financial impact. There is a substantial difference between a stockbroker who cannot place a trade and a marketing person who may be inconvenienced, but can work on other tasks. The critical step with quantifying the impact of poor network productivity is providing a realistic view of the impact of application and network downtime and degradation.

Employees require access to business applications

In this scenario, there are many variables tied to the calculation of the savings. This enterprise has 2,500 end users with an average weighted salary of $37,500 ($25,000/year + 50% benefit cost). Historically, the enterprise determines application and network degradation impact performance substantially approximately 2% of the time. Of the 2,500 users, approximately 10% will be negatively impacted for business critical activities (logistics and production are impacted, but sales, marketing, finance, HR, etc. are not affected).

With enhanced visibility and troubleshooting of network and application performance management, the amount of downtime is estimated to be reduced 45%. By reducing application downtime, the enterprise would save $84,375 in network user productivity (2,500 users x $37,500 loaded salary x 2% downtime x 10% end user impact x 45% savings).

Calculating a payback period

When deploying a new solution, enterprises today typically look at a payback period (when the savings pays for the solution's savings pay for the cost of deployment) or return on investment (ROI). Once you quantify the savings, the payback calculation is extremely simple. You take the total cost of deployment and divide that by monthly savings. The result is the number of months for payback (or payback period). The best way to do this is walk through an example of a typical enterprise customer deploying Visual Performance Manager.

In the scenario below, this organization is attempting to quantify the savings associated with deploying Visual Performance Manager. In the first section, they walk through the areas of savings as described earlier and determined the estimated annual savings would be $372,534 (please note each organization will have varying results for this analysis depending on the enterprise, applications, etc.). Once the estimated savings are quantified, the next step is to calculate the cost of the application and network performance management system, Visual Performance Manager. For this scenario, the complete cost for the initial deployment is $215,055.

Once you have these figures, you can calculate a payback period and a return on investment (ROI). The payback period determines the amount of time it takes for the performance management system to pay for itself. A payback is calculated by dividing the cost of deployment by the savings and multiplying by 12 months. In this scenario, the payback period is 6.93 months ($215,055 cost of deployment/$372,534 savings X 12 months).

A ROI quantifies how much return you will receive on an investment over a period of time. Most ROIs are calculated in 1-, 2- or 3-year projections. The ROI is calculated by dividing the total yearly savings by the cost of deployment. In this scenario, the 1-year ROI is 173 percent ($372,534 savings/$215,055 cost of deployment). When calculating a multi-year ROI, you multiply the savings by the number of years. The cost of deployment would include the one-time equipment cost and any annual costs such as maintenance. The 2-year ROI is almost 300% ($372,534 savings x 2 years/$247,860 – cost of hardware, software and 2 years maintenance). The 3-year ROI improves to 397 percent.

Payback/ROI Worksheet	
Savings with improved application and network performance	
Increasing revenue and production capabilities	$54,245
Optimizing bandwidth expenditures	$74,450
Decreasing number of overlapping tools	$22,785
Improving identifications and repair of virus and worm attacks	$65,550
Reducing application MTTR	$97,481
Lowering number of trouble tickets	$7,845
Reducing network configuration time	$15,678
Improving end user productivity	$34,500
Total savings per year	**$372,534**
Cost of Visual Performance Manager deployment	
Hardware	$44,750
Software	$137,500
Maintenance	$32,805
Total	**$215,055**
Payback period 6.92 months	
1-year ROI	173.23%
2-year ROI	299.69%
3-year ROI	397.14%

Armed with this qualitative analysis, your enterprise can make an informed decision on the benefits and savings associated with a proactive, application and network performance management system. You can run different scenarios using the same toolset to be more aggressive or more conservative and see how the payback period and ROI change.

Conclusion

Believing there are benefits to improved network and application performance is not enough for most organizations. The ability to quantify cost savings, improved productivity or reduced risks is a critical component in justifying an investment in application and network performance. This paper has listed eight key areas where cost savings can be quantified. Each organization will have different results and savings – some savings might be spread out evenly while others will be skewed to only one or two criteria.

Use the quantifiable results to calculate a payback period. Historically, Visual Performance Manager users have seen a payback period in the 5 to 7 month range. Contact your Fluke Networks representative if you'd like to walk through how quickly Visual Performance Manager pays for itself with improved performance for both the network and applications across the infrastructure.

About Fluke Networks

Fluke Networks is a leading provider of network and application performance management solutions. The company's technologies enable enterprises to reliably and securely manage the delivery of mission-critical applications across their infrastructure. Fluke Networks' products increase application and network availability, optimize the use of bandwidth, and reduce operating costs across traditional and IP-based infrastructures. For more information, visit **www.flukenetworks.com/vpm**.

Contact Fluke Networks: Phone 800-283-5853 (US/Canada) or 425-446-4519 (other locations).
Email: info@flukenetworks.com.