Netcordia™
Simply a better way to manage

Network Downtime, the Configuration Errors

Introduction

There is paradox at the heart of modern IT. Businesses invest in computer networks to increase agility, to extend their reach and speed of reaction, and to improve efficiency and reduce costs. The technology has been so successful in meeting all these objectives that no modern business can function without it.

But when it comes to managing these networks, it is a different story. Most businesses fail to fully employ the automated management tools available to them, and, in doing so, they fail to properly apply the proven techniques and expertise that would make their networks so much more reliable.

Paradoxically, most mission critical IT networks are configured and managed with little help from automated systems, but by humans working with few tools. They are sometimes poorly trained, very often lacking up to date information, and almost always working under pressure.

The result: the networks that drive modern businesses today are both a strength and a weakness. They often fail, at great expense, not because of underlying equipment problems, but because of human error in setting them up and running them. Modern business suffers from a serious, endemic and expensive problem that is not only largely avoidable, but is self-inflicted.

The Cost of Network Downtime

All users of corporate computers know about downtime. It is when the system, often for arcane technical reasons, either shuts down completely, or slows down to a point where it affects productivity. The causes range from hardware failure, software failure (bugs), incompatibilities between interlinked systems, and, very often, human error in setting up and managing the systems.

Many studies have been carried out to estimate the prevalence and cost of downtime, but few distinguish clearly between the role of the different servers and the networking equipment that links them together. There is, however, a consensus that networking problems, the focus of this paper, are both a major cause of downtime, and tend to create the most business damage.

Most businesses fail to fully employ the automated management tools available to them, and, in doing so, they fail to properly apply the proven techniques and expertise that would make their networks so much more reliable.

Although modern networks are more reliable than those of the 1990s and before, failures tend to have consequences that go far beyond those of only a few years ago, when IT's role was much less important. Modern IT operates in real time or near real time, systems typically serve far greater numbers of people, and there is far greater dependency on the systems remaining available. All of these factors mean that simple problems can rapidly create a business crisis.

The cost of enterprise and network downtime has been established by many studies. For example:

- **Infonetics Research (2004)** found that downtime of all enterprise applications cost, on average, 3.6% of annual revenues, at the 80 largest companies in the US.
- **Infonetics Research (2006)** found that businesses with 100-1000 employees lost 1% of their annual revenues in network downtime, or $867,000.
- **Gartner (2004)** put the average hourly cost of network downtime at $42,000. The average company suffered 87 hours of network downtime, amounting to $3.65 million.
- In a study which focused on application rather than network downtime, **Alinean (2004)**, a specialist in IT return on investment calculations, estimated that the loss of a supply chain application cost big companies $11,000 per minute, an e-commerce application $10,000 per minute, and a customer service application $3,700 per minute.

All of these numbers, as the authors readily acknowledge, need to be treated carefully, because the calculations are necessarily complex and involve intangible factors. The cost of downtime needs to take into account several factors, ranging from the cost of repair (mostly labor), the cost of lost productivity by end user staff, and above all, any loss of business or goodwill that may result. Some web sites have calculators to enable individual organizations to calculate the impact of downtime.

Naturally, these figures are difficult to quantify, and vary widely according the size and type of company, and from industry to industry.

A commodities trader or major e-commerce site is likely to be rather more seriously affected than a company producing, say, educational books.

The cost of a web site outage at Amazon, for example, has been independently calculated at $350,000 an hour. The company suffered a two hour outage in 2006, and several during the peak shopping season in 2001.

The true impact of network downtime can sometimes be difficult to spot. Although network downtime can be sudden and absolute, it is just as likely to occur partially or gradually, so that users don't notice until business has already suffered. Performance of key applications can gradually deteriorate, so that users or even customers become increasingly frustrated and may stop using the system long before it has officially "failed". Sometimes some departments or applications will be affected, but others will not.

A large study by Vanco, a managed network operator, found that most network downtime was caused by configuration errors or by users trying to resolve problems beyond their technical competence.

Pilot Error

Perhaps the most remarkable aspect of the downtime problem is that one factor is consistently and repeatedly identified as a primary (if not the primary) cause: human error. This applies to application development, systems management and, of course, network management.

Configuration problems are a particular problem. A large study by Vanco, a managed network operator, found that most network downtime was caused by configuration errors or by users trying to resolve problems beyond their technical competence.

The authors of a detailed study into network failures by the University of Michigan found that 59% of the problems causing downtime in IP networks pertained to routing management, and more than a third of these problems were directly caused by configuration errors.

In the wake of a much publicized network problem at Microsoft in early 2001 (see below), Gartner calculated that 80% of all downtime was caused by "mismanagement and misconfiguration".

Certainly, there have been many examples where configuration errors have caused expensive problems at major organizations. For example:

- Just a day after the software giant had announced a $200 million advertising campaign promoting the notion of the IT powered "Agile Business", all of Microsoft's major web sites were down for up to 24 hours. The company suspected it had been the victim of a denial of service attack – but the cause was a simple router misconfiguration.
- A large medical research organization discovered, on installing Netcordia's NetMRI appliance, that over 1,000 network devices were wrongly configured, causing the network to run much more slowly than it should for most users. Most were then quickly fixed.
- An engineer at a large financial institution made a configuration change, which brought the stock trading system to a halt. The company lost so much money that it banned any network changes during trading hours.
- A network administrator at a major financial organization accidentally made some changes to the Active Directory. As a result, the trading desk was shut down for eight hours, costing the company

millions of dollars. A simple restore of the original settings solved the problem.

Why do trained staff make so many apparently simple mistakes? The answer is largely one of complexity and workload. Although most engineers understand the problems, they are often overwhelmed by the amount and type of equipment, the amount of alerting information that these systems produce, and by the large number of dependencies that exist between devices. Often, there is insufficient time to properly check configurations or carry out preventative maintenance.

Solutions

If network downtime is so expensive, and if so many of the problems are caused by apparently simple errors, then why is it so still so common? And what can managers do to reduce the level and impact of downtime at their organizations?

Certainly, most major organizations, especially in areas such as financial services and e-commerce, have attempted to reduce network downtime by investing heavily in expertise, in tools and in network and systems redundancy. Many suppliers, too, have developed products aimed at improving network management and thereby reducing downtime.

The most effective organizations are able to learn both from their mistakes and those of others, and to distribute information quickly.

Experience has shown, however, that as networks continue to become more complex and interdependent, configuration and management problems continue to cause expensive downtime, even when sophisticated network management products are in place.

In recent years, some experts have been advocating a more rigorous, business level focus on the problem, sometimes with dramatic results. Among the areas they focus on are:

– **Policy**. The most effective organizations are those that have clear policies for how they tackle complex tasks, and this applies to configuring and managing networks. For example, there should be policies on which traffic is allowed through the firewall, on which applications are given bandwidth priority, on how new user accounts are set up and managed. There should be clear policies for dealing with under performing equipment, and on how failures should be dealt with. There should also be strict policies governing who is allowed to change network settings, and how they should be configured. The network management company Netcordia maintains a database of best practices and policies, and its NetMRI appliance enables business to maintain and develop this.

– **Automation**. They may be well intended, but humans make mistakes, frequently miss minor details, or occasionally act outside or beyond their competence. Automation should be used wherever possible, both in adopting and enforcing policies and in configuring equipment.

The most effective organizations are able to learn both from their mistakes and those of others, and to distribute information quickly. Again, the network management company Netcordia uses its rules-based database that administrators can use to configure thousands of network devices directly, ensuring the most stable and effective performance. Any attempts to change the configurations are instantly detected using the configuration management database, which automatically collects data from all the devices on the network.

– **Simplification**. Modern networks are managed with the help of a proliferation of tools, among them IBM's Netcool, Cisco's CiscoWorks and Hewlett Packard's OpenView.

These tools are widely used, and often play a key role in minimizing network downtime. But they are mainly used to identify problems and their causes after they have occurred – not before.

These tools can be difficult to use and provide a variety of alerts. Often, a failure in one device triggers further problems downstream, creating a proliferation of alerts, requiring the use of further analysis to discover the original problem.

The use of a single tool, or a few simple, integrated tools, for configuration and performance management is likely to help. This will enable better network configuration and management in the first place, so that fewer errors are made, and problems are spotted and resolved when they do occur.

– **Prediction**. Perhaps nothing is more important in reducing network downtime than in preventing it in the first place. Netcordia's NetMRI is particularly innovative, adopting techniques from the world of analytics and business intelligence, where centralised databases and easy to use tools enable managers to quickly understand complex business situations.

The analytics engine draws not only on current network information, but on information from manufacturers and on Netcordia's database of industry best practice. This enables Netcordia to provide an instant, BI style "scorecard" of the health of the network at any point, and, if it is deteriorating, to drill into the reasons why. Evidence from Netcordia suggests that the use of its analytics approach can dramatically reduce network downtime and improve the time to repair.

- **Communicating the payback**. When it comes to reducing network downtime, the most effective IT and network managers will be those that are best able to communicate the need to solve the problem to their finance director or the board. This is, fortunately, not a difficult challenge; the return on investment on effective network management tools, such as NetMRI, is clearly apparent – in terms of reducing labour costs, maintaining productivity, and avoiding the danger of losing customers, revenue and goodwill.

Networked World

Will network downtime get better or worse in the coming years? Over time, improvements in ease of configuration, auto-detection of problems, and better equipment reliability should reduce the number and the impact of network problems.

But a number of factors are also making downtime more likely.

Those organizations that invest in a solution that captures, distributes and applies best practices in configuration and network management are far more likely to reduce or eliminate the problem of network downtime.

Increasing equipment and network architecture complexity, coupled with the introduction of demanding new services such as voice over IP and video, are likely to drive up the likelihood of configuration problems occurring.

These new applications such as voice, in particular, are more likely to suffer from low latency problems, meaning that network degradation, as opposed to outright outages, will have a greater impact. The greater spread of applications, and of the network infrastructure, over many geographies, is also likely to lead to more problems.

Those organizations that invest in a solution that captures, distributes and applies best practices in configuration and network management are far more likely to reduce or eliminate the problem of network downtime.

This whitepaper was written on behalf of Netcordia. Netcordia offers enterprises, service providers and government agencies simply a better way to manage their networks. Able to deploy in less than 30 minutes, Netcordia's NetMRI provides the continuous analysis, reporting, and troubleshooting capabilities necessary to manage multivendor networks. This automated solution provides visibility into the network to easily and quickly identify and track issues related to compliance, configuration, performance, security, and IP Telephony which can lead to network instability and service interruption.

Netcordia has been recognized as one of the top companies to watch in network management by Network Computing and one of the Fierce15 top emerging companies in IP Telephony by FierceVoIP. **(www.netcordia.com)**

For more information on Netcordia and our full range
of automated network management solutions,
please call **410-266-6161** or **visit netcordia.com**

Netcordia

Simply a better way to manage