

# A Best Practice Guide to Network Troubleshooting

## Introduction

**Network Troubleshooting (or Problem Management in the world of ITIL<sup>1</sup>) is an art, which when properly applied can result in cost savings, reduced downtime and increased network stability. Applying set processes and understanding the obstacles to success can supply the Network Manager with quick wins, improve customer and user perceptions of the service being provided, and build bridges between the various support functions delivering IT for the Business.**

Based upon Kenson's wealth of experience in the resolution of network centric issues and our adopted Best Practice methodologies, this document is designed to provide an understanding of the political obstacles to troubleshooting and how they may be overcome, whilst giving an outline of the Troubleshooting process and ways it can be enhanced to be more efficient.

Built up over many years, the arsenal of cost effective and feature rich tools used by Kenson in the delivery of their consultancy based services is provided as an example of products, which may be used to aid the Troubleshooter in the resolution of their issues.

The introduction of the controls and processes described in this document will be of benefit to all parties – the Business and its Customers, the IT function supporting the Business and, of course, the users of the network.

**Kenson Network  
Engineering Ltd**  
Corinium House  
Units 101/102  
Cirencester Business Park  
Love Lane  
Cirencester  
Gloucestershire  
GL7 1XD

**T:** +44 (0)1285 647900  
**F:** +44 (0)1285 643686  
**E:** enquiries@kenson.co.uk  
**W:** www.nut.eu.com/kenson

 **Nu Technologies**

Kenson are part of the  
Nu Technologies Group

## The Need to Know Basics

The goal of Network Troubleshooting is to identify the root cause of a service effecting issue so that a permanent fix or workaround can be established and applied and the effected network service can be restored in as short a time as possible.

There are a number of basic principles applying to all Network Troubleshooting situations, that people embarking on the Troubleshooting process need to know.

The first unwritten law of Network Troubleshooting is that the issue has surfaced for one of only two reasons; something has failed or something has changed.

Although tools are available to detect component failure and despite changes to the network infrastructure often being under the control of a Change Management process, the single most inhibitive factor in the timely resolution of network issues is the lack of full, controlled and accessible documentation of the devices which make up the network service and the relationships between them. Documentation, and more importantly, up to date and accurate documentation of the network, its hardware and its configuration is crucial for the control of the IT infrastructure. Without it, delays and mistakes are unavoidable.

In ITIL<sup>1</sup> terms, all information relating to the network infrastructure is stored in a documentation store known as the Configuration Management Database (CMDB). The information held in the CMDB is kept up to date by strict adherence to Change Management. Without a comprehensive Change Management process, the information stored in the CMDB quickly becomes out of date and inaccurate, a mere snapshot of a

previous moment in time. To overcome this, it is recommended that the IT infrastructure is audited periodically to verify CMDB integrity and all changes to the network are documented by Change Management by use of the Forward Schedule of Change otherwise known as the Change Log.

Change Management and the inherent controls it imposes will always allow regression to the previous, stable state of being.

Whilst on the subject of Change, it is important to note that changes carried out to test the resolution of a network issue should be applied one at a time. A number of changes applied simultaneously to 'blitz' the issue will be hard to regress and may introduce a secondary problem, for which the root cause then becomes difficult to trace.

Control of the changes applied and the documentation of all that is done cannot be emphasised strongly enough.

A major benefit to this approach is the removal of assumption from the view of the Troubleshooter. Assumption feeds the Troubleshooter red herrings and is to be avoided if efficient and timely progress is to be made.

The Troubleshooter must also bear in mind that, although the information gathered regarding the symptoms of the service effecting issue is primarily subjective, i.e. based on the user experience, interview and, to a certain extent hearsay, theories and hypotheses devised for the resolution of major issues must be based on fact, and, overall, be objective.

Once the root cause of a problem has been established it is important that details of the issue, its symptoms, verification techniques, the cause and the fix for the problem are documented and added to an extensible Knowledge Base controlled by the Configuration Manager so that the service effecting issue can be readily identified and the fix applied should it re-surface in the future.

### **The Politics of Troubleshooting**

Often the reason for independent consultants such as Kenson being employed for troubleshooting of persistent network issues is to circumvent the artificial boundaries formed by the segregation of the internal IT Support functions. The focus becomes the apportioning of blame as the Server Team blames the Network Team, which in turn blames the Build Team. This situation is always counter productive. Proof of responsibility is required before issues can be owned and ultimately resolved. The role of Troubleshooter or Problem Manager in ITIL<sup>1</sup> terms, should transcend the boundaries of individual IT Support teams and focus objectively on the evidence presented to him. Initially, his role is to hear the problem, preferably observe the problem, to define and describe the problem without prejudice or assumption so that all parties can discuss ways to resolution without feeling threatened or undermined.

It is vitally important for all IT Support personnel to recognise that their primary role is to support the Business in the delivery of IT Services and all groups should work together to this end.

In the event that internal conflict exists, it is often a time saving and economically justifiable option for an external resource to be employed to bridge the political divide so long as internal processes are adhered to and knowledge and thereby control is returned to the internal IT Support function once the root cause of the service effecting issue is identified or resolved.

### **The Troubleshooter's Toolkit**

There are a large number of tools available to the Network Troubleshooter, which will assist in the diagnosis, testing and resolution of network issues. The following list of tools is not exhaustive, but gives a good indication of the tools, which are invaluable to the Troubleshooter when presented with network issues.

Simple Network Management Protocol (SNMP) based network monitoring systems such as SolarWinds Orion or HP OpenView will provide reporting on availability, capacity, utilisation, error conditions and other key metrics for active, manageable network devices. Whole or partial network views can be provided with colour coding based on preset thresholds. Other features will be alerting via e-mail or SMS and historical reporting for trend analysis. Many SNMP Management Systems, including those listed above, have an in-built Web Server, which can allow read-only access to all Support

functions via a standard Web Browser and allow views to be tailored to the needs of individual groups.

In order for devices to be included in the scope of SNMP Management System coverage, devices must be at least IP addressable, so that they can respond to ICMP ping packets for Availability statistics. Full management requires that the devices have an operational, standards compliant SNMP Agent, which will forward statistics such as Interface, CPU and Memory Utilisation when requested by the Management System. This should be considered when equipment is purchased as cheaper devices may not have this facility and fault diagnosis will be more difficult as a result.

Although no longer the universal panacea, which was the case in the days of Shared Ethernet, a protocol analyser such as Network General's Sniffer or the freely available open source Ethereal, is still of considerable use in today's packet switched infrastructure for the diagnosis of application based issues. Conversations and packet flows can be captured through a mirrored switch port for real time analysis or saved for later interpretation. General Broadcast and Multicast traffic trends can also be identified and their impact assessed.

It is often the case that Wide Area Network (WAN) connectivity is provided via a managed service, which does not allow analysis of conversations or capacity relative to business critical applications, or by a Virtual Private Network (VPN) across the Internet. In cases like these the use of a traffic analysis and packet shaping system such as Allot's NetEnforcer or Packeteer's PacketShaper may be an appropriate tool to guarantee delivery of critical data by means of prioritisation.

In a Cisco routed environment primarily, although not exclusively, a wealth of information with regard to conversations and protocol distribution may be obtained from the Cisco proprietary Net Flow feature. Tools such as Crannog NetFlow Tracker provide a storage location for NetFlow information and a Graphical Interface, which allows the creation of easily interpreted reports of historical NetFlow data. It is worth noting that Cisco NetFlow v.9 is proposed as the basis for a standard IP Flow Export Protocol by the IETF IPFIX working group.

As a result, it is expected that NetFlow Tracker's coverage will extend to most network equipment suppliers in the near future.

Data stored in the CMDB relating to the network or services supported by the network will often be unwieldy, disparate and difficult to convey to other support functions at the problem diagnosis stage. Use of a data visualisation tool such as netViz enables information drawn from any ODBC compliant database to be presented in a simple to understand, hierarchical view, which will provide quick wins for all parties involved. Publishing of the pictorial representation provided by netViz to a Web Server (netViz WebView) will once again provide universal read only visibility of service interdependencies to all support functions. More collaborative systems such as netViz Enterprise Server can be integrated to Change Management systems, providing work orders, test plans and maintaining synchronous visualisation and back end data.

When considering the procurement of tools for network management and troubleshooting in particular, it is important that the cost of training is factored in to the budgetary costs. Tools are of optimal use only if the personnel using them can obtain the maximum benefit from them. All manufacturers and some resellers, such as Kenson, provide comprehensive training, which will facilitate full exploitation of the supplied tool's features.

Three outputs from Best Practice for IT Service Support<sup>2</sup>, previously mentioned in this article, are of invaluable use in the identification of network issues. These are the CMDB and the Knowledge Base from Configuration Management and the Forward Schedule of Change from Change Management. These outputs will allow up to date understanding of the network infrastructure and relationships between devices as well as the timely identification of issues, which have re-occurred.

Finally, it is useful to note that the Internet is a rich source of information once an issue has been correctly diagnosed. Most issues or network problems have already been encountered by someone else. Most equipment manufacturers and software providers have publicly accessible Knowledge Bases. Forums dealing with a wide range of issues are available for searching and possible resolutions are often posted on them. However, caution should be exercised as these solutions are rarely definitive and testing is always required to assess their validity.

## Process and Methodology

As previously mentioned, the goal of Network Troubleshooting is to identify the root cause of Service effecting issues so that a permanent fix or workaround can be devised and applied.

In order to do this, it is necessary first of all to gather as much information as possible regarding the issue itself, the circumstances whereby it manifests itself and the boundaries of the issue.

Interviews should be held with those responsible for the devices and systems adversely effected by or potentially contributing to the service failure or performance issue as well as those users impacted by the issue.

Information gained from interviews will include the impact of the issue (systems effected, geographical scope of the issue, hardware impacted, software impacted, time of impact etc), the pattern of the issue (time of occurrence, sequence of events leading to the incident etc) and the history of the issue (when first reported, any previous attempts at resolution etc).

If the Troubleshooter is an external resource such as an independent consultant, it will also be necessary to interview the IT Manager of the business so that the rules of engagement and the incumbent internal processes can be clearly understood by both parties.

The Troubleshooter should deal in facts but bear hearsay in mind, with the collated information providing a subjective view of the issue.

As a result of the information gathering exercise, it should be possible to observe the issue first hand, describe the issue without prejudice or assumption and define the issue in objective terms. By understanding the issue and the circumstances leading to its occurrence, it should be possible to replicate the issue and provide a baseline against which testing and authorised changes can be implemented. A clear understanding of the circumstances whereby the issue is triggered may also lead to an understanding of the underlying cause of the issue.

Once the issue is understood and documented, this information should be shared with all teams supporting the effected service so that agreement may be reached that it represents an accurate assessment of the Problem and that there is general consensus as to the boundaries of its impact. The sharing of objective information in this way may also lead to root cause identification as the individuals within the support teams are no longer defensive in attitude, but part of a cohesive effort to resolve the issue defined.

Data from the CMDB should be examined to understand the devices comprising the effected service, their interrelationships and dependancies.

Information such as the timing of the initial reporting of the incident may be used to compare with changes made to the network described in the Forward Schedule of Change and may identify the root cause of the issue as being a by-product of a previous change.

The definition of the issue should be used to consult internal and external Knowledge Bases regarding previous incidence of the issue. This can provide quick wins by identifying a known fix or workaround and reducing the time to resolution.

If the issue has not been previously documented or there is no known fix, further diagnosis or testing may be required before a fix can be devised. The Troubleshooter should understand any adverse effect fault diagnosis or disruptive testing will have prior to raising a Request for Change (RFC), which will sanction the testing process. The testing process should be documented prior to implementation so that any gaps may be highlighted and impact on other systems identified. Once again, collaboration with support teams at this stage will prove beneficial.

Once ratified and scheduled by Change Management, testing should be carried out under strictly controlled conditions and the outcome documented. Again, this will highlight any deficiencies in the testing process prior to subsequent action plans.

Testing should be honed and redefined in a logical sequential manner until identification of the root cause of the issue has been accomplished, at which stage, a fix or workaround can be devised and passed to Change Management for implementation.

Changes at this stage should be carried out one at a time so that the effect of the change can be properly understood. Use of a documented implementation plan is vital as this will allow regression if the fix proves to be ineffectual. If a change is beneficial but not

effective in the resolution of the issue at hand, this can be documented and scheduled for later implementation once the primary issue has been resolved, but should be regressed so that its effects do not add confusion to the diagnosis of an effective fix.

When the fix has been applied, an independent tester, such as an effected user should be employed to assess the validity of the fix using a documented test plan. If the testing indicates that the issue has been resolved, a bedding in or monitoring period should be set so that the impact of the fix on the network as a whole may be assessed.

Finally, the symptoms of the issue, verification techniques and details of the fix or workaround should be logged in the internal Knowledge Base so that it can be readily identified and the fix applied in a timely manner should the issue resurface at any stage in the future.

### **Proactive Steps**

Excessive Service Outages are often the result of the existence of single points of failure in the infrastructure delivering the Service. In an ideal world, these should be identified and steps taken to build a degree of redundancy into the infrastructure so that the impact of network issues on Business Critical Systems can be minimalised. Unfortunately, resilient solutions are expensive and difficult to justify as the redundant equipment is largely unused and when it is called into play, it is often unnoticed by Senior Management.

In order to convince Senior Management of the benefits of bolstering the network infrastructure to include resilience for Business Critical systems, it is necessary to provide some cost justification.

The cost of downtime to a large enterprise business is estimated by Gartner to be approximately \$42000 per hour. Even in a small organisation it is estimated to cost 1% of annual revenue, which can equate to a substantial sum.

To take a more scientific approach, the annual cost to your business of network downtime and therefore the cost justification for the implementation of increased resilience within the network infrastructure can be calculated using the following equation:

$$\text{Annual Cost of Network Downtime} = h + k + n + r$$

Where:

- a = Annual Revenue of Company
- b = Number of Employees
- c = Annual Revenue per Employee (a / b)
- d = Annual Total Hours per Employee
- e = Average Revenue per Employee per Hour (c / d)
- f = Total Man Hours to Restore Service
- g = per Hour Cost of Restoration Services
- h = Cost of Labour to Restore Service (f x g)
- i = Total Hours Service is Down
- j = % of Employees Unproductive during downtime
- k = Cost of Employee Downtime ((b x e) x i x j)
- l = Number of Sales per Year
- m = Estimated Number of Sales Lost Due to Outage
- n = Cost of Lost Sales Opportunity ((a / l) x m)
- o = Total Number of Customers Last Year
- p = Average Revenue per Customer (a / o)
- q = Number of Customers Lost due to Service Failure
- r = Cost of Loss of Customers and Reputation (p x q)

By reducing downtime, the Business benefits from increased Customer and User satisfaction, increased productivity and increased revenue. However, the reduction in network downtime resulting from cost justified resilience solutions must be reported to Senior Management so that successes can be recognised and Return on Investment (RoI) can be identified and confirmed.

## Company Overview

Since its inception in 1989, Kenson's primary mission has been the provision of best practice professional services and best of breed products in network management and operations. Working in partnership with our customers, we work to improve network management techniques, increase network reliability and efficiency and provide the skills necessary to meet the challenges of today's business critical networks.

Kenson is part of the Nu Technologies group of companies.

Further information relating to Kenson's Products and Services can be obtained by access to our web site at <http://www.kenson.co.uk>, by e-mail to [enquiries@kenson.co.uk](mailto:enquiries@kenson.co.uk) or by telephone on +44(0)1285 647900.

- 1 ITIL is the Information Technology Infrastructure Library and ITIL® is a Registered Trade Mark and a Community Trade Mark of the Office of Government Commerce (OGC). ITIL provides a customisable framework of Best Practice methodologies for the Management and Delivery of IT Services.
- 2 Best Practice for Service Support is published by the Office of Government Commerce (OGC) and is part of the IT Infrastructure Library.

### **Kenson Network Engineering Ltd**

Corinium House  
Units 101/102  
Cirencester Business Park  
Love Lane  
Cirencester  
Gloucestershire  
GL7 1XD

**T:** +44 (0)1285 647900  
**F:** +44 (0)1285 643686  
**E:** [enquiries@kenson.co.uk](mailto:enquiries@kenson.co.uk)  
**W:** [www.nut.eu.com/kenson](http://www.nut.eu.com/kenson)

 **Nu Technologies**

Kenson are part of the  
Nu Technologies Group