

Performance Management Best Practices for Broadband Service Providers

Contents

- 1 Introduction
- 2 Scope
- 3 Performance Management Terminology
 - 3.1 Performance Metrics
 - 3.2 Key Performance Indicators
 - 3.3 Service-Level Agreements
 - 3.4 Operation-Level Agreements
- 4 Typical Issues and Problems Affecting Performance
 - 4.1 Performance Management Building Blocks
- 5 Performance Management Operation Workflow
- 6 Implementing Performance Management
 - 6.1 Define Performance Goals
 - 6.2 Identify the Performance Resource Constraints
 - 6.3 Define Performance Metrics
 - 6.4 Document Business Objectives of Performance Management
 - 6.5 Document Performance Management SLAs
 - 6.6 Document Performance Management OLAs
 - 6.7 Performance Data Collection
 - 6.8 Performance Data Reporting
 - 6.9 Measure Network Performance
 - 6.10 Perform a Proactive Fault Analysis
 - 6.11 Performance Baselining
 - 6.12 Review the Performance Baseline
 - 6.13 Document a Risk Analysis Methodology
- 7 Methods to Refine Performance Management
- 8 Broadband Case Study

- 8.1 Business Requirements
- 8.2 Typical Architecture of a Fixed Broadband Service Provider Offering
- 8.3 Data Collection Sources and Points
- 9 References
- 10 Conclusion

1 Introduction

The primary function of an ideal performance management system is to optimize the use of the network and applications so as to provide a consistent and predictable level of service. Once this goal is established, then the focus of performance management is to optimize the performance of the network and applications in order to comply with service-level agreements.

Performance management involves the following:

- Configuring data-collection methods and network testing
- Collecting performance data
- Optimizing network service response time
- Proactive management and reporting
- Managing the consistency and quality of network services

Performance management is the measurement of network and application traffic for the purpose of providing a consistent and predictable level of service at a given instance and across a defined period of time. Performance management involves monitoring the network, application, and service activity and adjusting designs and configurations in order to meet performance requirements or improve performance and traffic management.

2 Scope

The scope of this paper is as follows:

- Set the foundation for understanding performance management principles and the parameters within service quality agreements, also called service-level agreements (SLAs).
- Define proactive measures to make network, application, and service performance visible to you and your organization.
- Provide guidance supported by examples and scenarios that you can apply in your network.

This paper discusses the following topics:

- Performance management terminology
- Typical issues and problems affecting performance
- Performance operation workflow
- Implementing performance management
- Broadband case study

This paper will use service provider broadband networks as examples, while the subjects are relevant for any network type.

3 Performance Management Terminology

In terms of networking, performance management is the configuration and measurement of network traffic for the purpose of providing a consistent and predictable level of service. Performance management involves monitoring network activity and analyzing and adjusting network design and configuration in order to improve performance and traffic management.

Performance monitoring involves the continuous collection of data concerning the performance of network elements. Performance monitoring is designed to measure the overall quality of performance, using monitored parameters to detect degradation.

Performance analysis: The collected performance records may require additional processing and analysis to evaluate each entity's performance level. Performance analysis includes the following functions:

- Creating recommendations for performance improvement
- Evaluating threshold policy
- Forecasting network usage
- Capacity analysis of the network

3.1 Performance Metrics

Performance metrics provide the mechanism by which an organization measures critical success factors. Performance metrics vary from business to business. They indicate how you will determine whether you have carried out the critical success factors that you identified, and indicate the kind of data you will need to gather.

Performance metrics are as follows:

- Connectivity (one-way)
- Delay (both round-trip and one-way)
- Packet loss (one-way)
- Jitter (one-way) or delay variation
- Service response time
- Measurable SLA metrics

3.2 Key Performance Indicators

A Key Performance Indicator (KPI) is a quantifiable metric of the performance of essential operations and/or processes in an organization. Hence, KPIs can do the following:

- Reflect the performance of network operators in achieving their goals and objectives.
- Reflect strategic value drivers rather than just measuring non-critical business activities and processes.
- Measure the health of the networks, servers, and applications and help ensure all individuals at all levels are using consistent strategies to achieve shared goals.
- Provide the focal point for standardization, collaboration, and coordination.

Consider the following during the KPI definition and development process:

- What should you measure? And how often should you measure?
- How complex should the metric be?
- What should you use as a benchmark?
- How do you ensure the metrics support your strategies?

3.3 Service-Level Agreements

An SLA defines and regulates other required performance management processes. Service-level management is a proven methodology that helps with resource issues by defining a deliverable and creating two-way accountability for a service tied to that deliverable.

Create an SLA between users and the network organization for a service that includes capacity and performance management. The service should include reports and recommendations to maintain service quality. However, the users must be prepared to fund the service and any required upgrades.

3.4 Operation-Level Agreements

Operation-level agreements (OLAs) are internal “back-to-back” agreements that define how two different organizations will work together to support the delivery of defined IT services to customers and users. While an OLA is very similar to an SLA, it has some key differences. An OLA defines how departments will work together to meet the service-level requirements (SLRs) documented in an SLA. If you do not have formal SLAs in place, you are still delivering IT services, and a Service Catalog will do instead.

4 Typical Issues and Problems Affecting Performance

Network operators are continually challenged to efficiently deliver business-critical applications while running the network at optimum performance. Also enterprises that deliver some form of service to their customers need an effective strategy for monitoring network conditions prior to potential issues becoming serious problems. Poor network performance can result in lost productivity, opportunity, and revenue, and increased operational costs. Hence, detailed monitoring and tracking of the network, applications, and users are essential in optimizing network performance.

Table 1 lists the problem categories that are addressed by performance management.

Table 1. Problem Categories and Performance Management Benefits

How Performance Management Can Help You (Benefits)	Problem Category
Informs the operator of impending performance deterioration	Application performance
Manages and prioritizes traffic	Capacity planning
Informs the operator of impending performance deterioration	Proactive fault management
Informs the operator of network availability and KPI breaches	Network outage impact; delay, jitter, and availability requirements; network latency and interface errors and discards
Communicates network performance in real time	CPU and memory utilization and network traffic in/out
Monitors bandwidth and categorizes network application traffic	Unbalanced bandwidth utilization and network hardware resource utilization
Provides the tools to pinpoint causes of performance deterioration or failure	Subscriber session performance

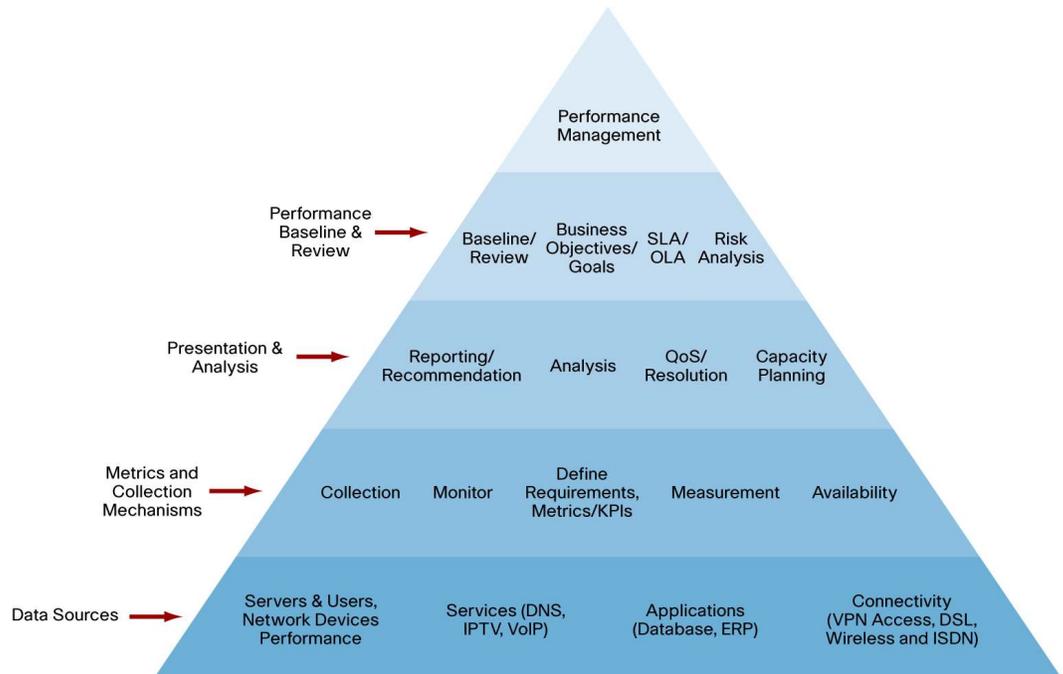
Performance problems are usually related to capacity. Applications are slower because bandwidth and data must wait in queues before being transmitted through the network. In voice applications, problems like delay and jitter directly affect the quality of the voice call.

4.1 Performance Management Building Blocks

Performance management provides functions to evaluate and report upon the behavior of network equipment and the effectiveness of the network.

It is essential for network operators to understand the goals and objectives of performance management—especially the direction in which their business shall go. Continuous and real-time reviews of network performance help to identify and eliminate problems before they impact their business services and customers. Figure 1 illustrates performance management building blocks

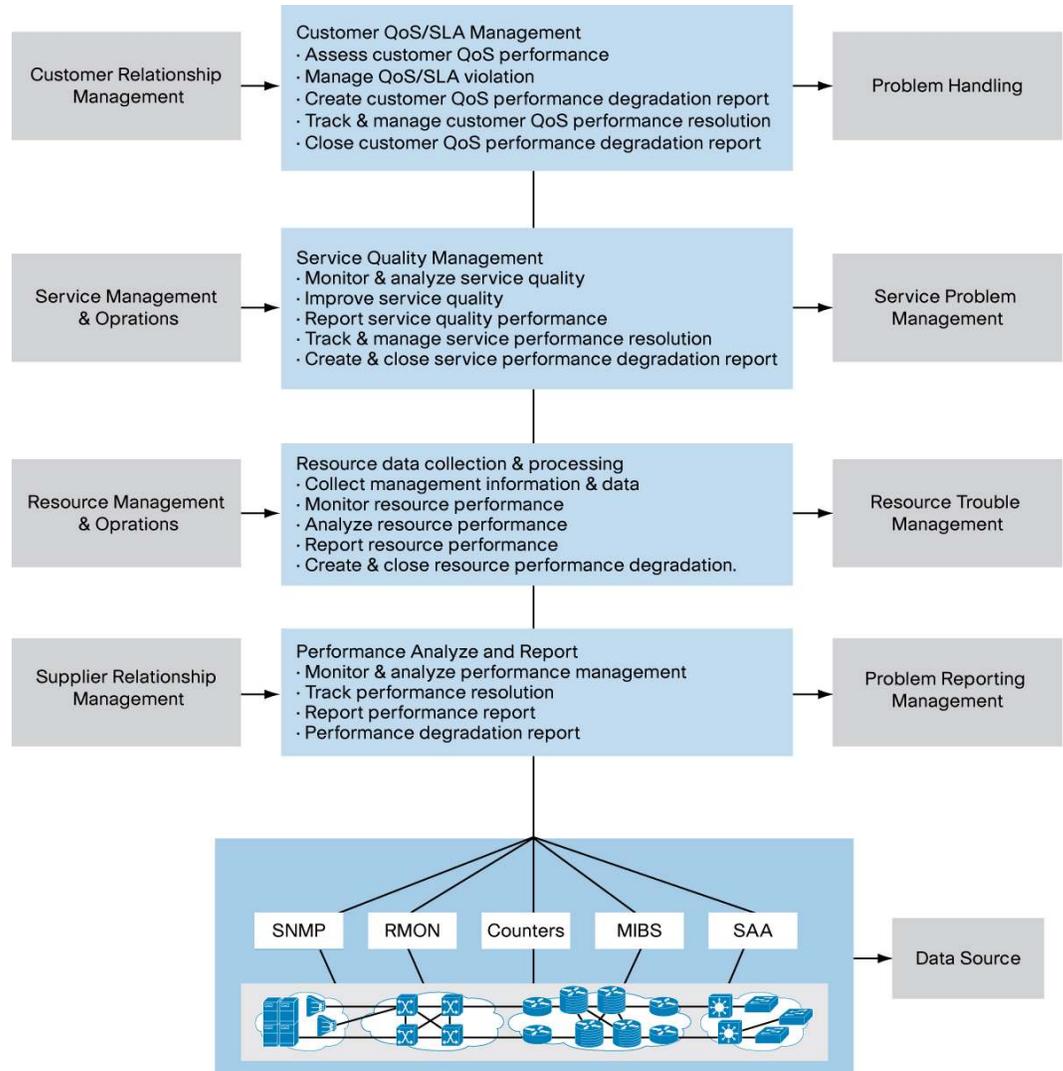
Figure 1. Performance Management Building Blocks



5 Performance Management Operation Workflow

Figure 2 illustrates a high-level operation workflow for performance management.

Figure 2. Performance Management Operation Workflow (Source: TMF NGOSS eToM)

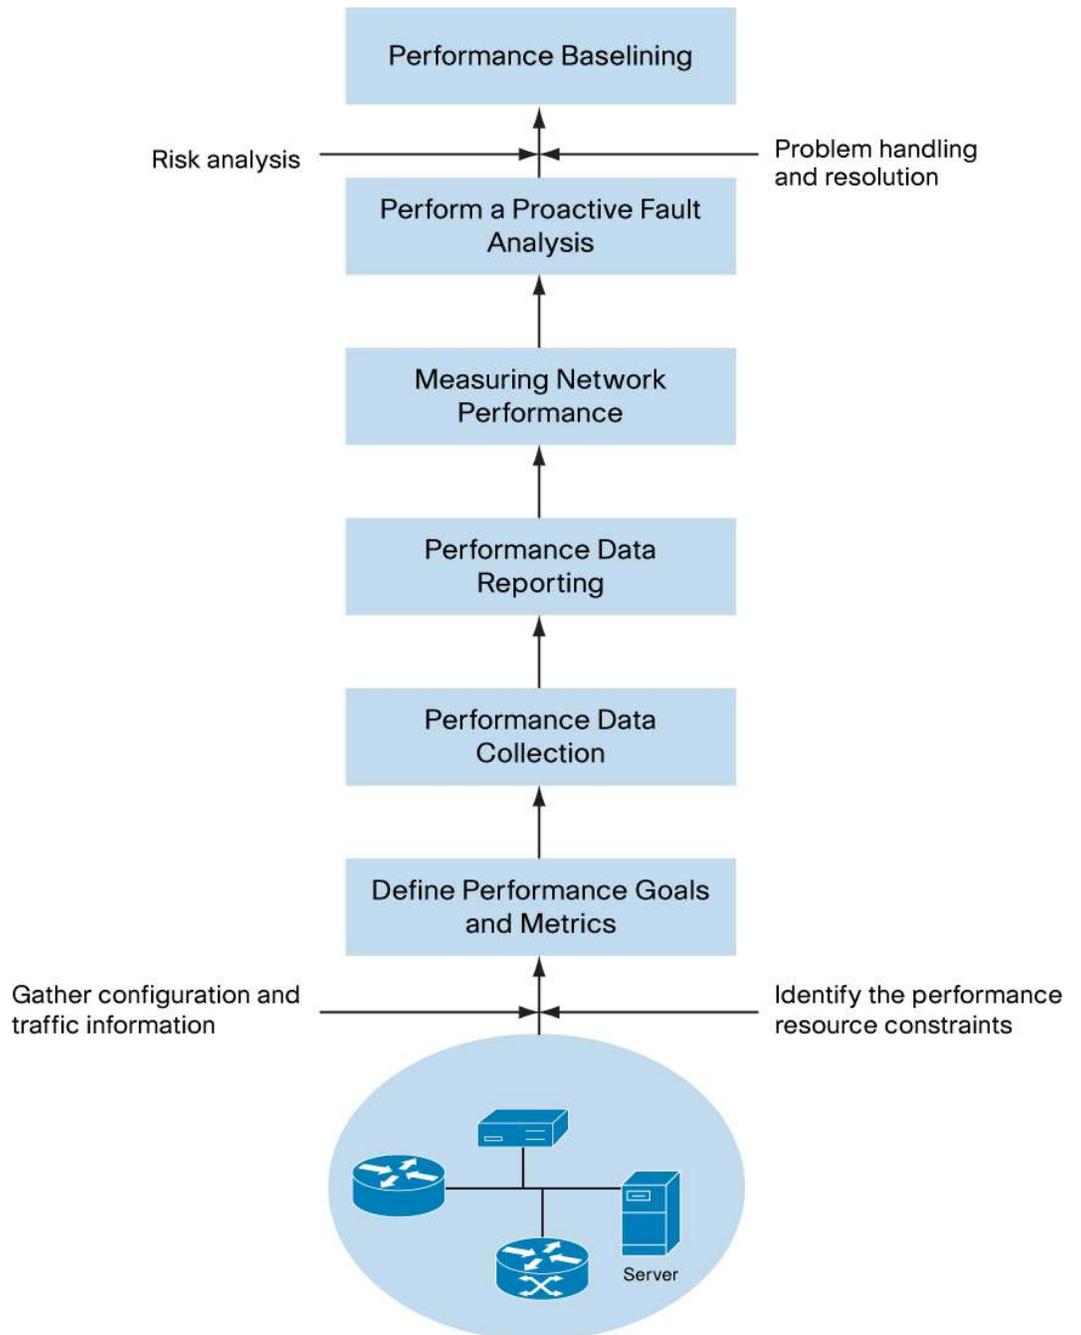


Defining the performance variables for a network provides a business foundation upon which you can build precise definitions of the features desired in your network. If you fail to develop an operational concept for network management, it can lead to a lack of goals or goals that constantly shift due to customer demands.

6 Implementing Performance Management

The first step is to define the required features and services before planning capacity, designing the network, and implementing performance management. This step requires that you understand networks, applications, basic traffic flows, user and site counts, and required network services.

Figure 3. Implementing Performance Management



6.1 Define Performance Goals

As part of the performance management process, it is essential to define the goals for the network, applications, and supported services in a way that all users can understand.

Each of the performance goals should be defined in a measurable way. Include the use of availability and response-time metrics tied into a system of notification when thresholds are exceeded.

- **Availability:** Availability is the measure of time for which a network system or application is available to a user. From a network perspective, availability represents the reliability of the individual components in a network.
- **Response time:** Response time is the best measure of customer network use and can help you gauge the effectiveness of your network.
- **Throughput:** Throughput is used to measure the overall data transfer capability of a network
- **Gather baseline data:** Perform a baseline review of the current network/system prior to a new solution deployment and after the deployment in order to measure expectations set for the new solution. This baseline helps determine if the solution meets performance and availability objectives and benchmark capacity.

To achieve an ideal network management system, you must actively integrate the components of performance management into the network/system.

6.2 Identify the Performance Resource Constraints

Capacity planning is the process by which you determine requirements for future network resources in order to provide the necessary performance or availability for business-critical applications. In this sense, capacity planning is preemptive performance management.

Performance management has the following resource constraints:

- Number of routers and switches
- Interfaces and virtual circuits
- Queuing, latency, and jitter
- Media capacity, speed, and bandwidth
- Application characteristics
- Device throughput
- Servers components
- Network services
- Connectivity media capacity

6.3 Define Performance Metrics

Performance metrics indicate how you will determine whether you have carried out the critical success factors that you identified and indicate the kind of data you will need to gather:

- Availability
- How often you met or missed SLA requirements
- Number of problems that occur
- Minimum, maximum, and average time to close a trouble ticket in a given priority

- Breakdown of problems by problem type (hardware, software crash, configuration, power, user error)
- Breakdown of time to resolve each problem type

6.4 Document Business Objectives of Performance Management

Business objectives should be documented in an in-depth study of how the enterprise or service provider wants to operate its business. The analysis focuses on existing infrastructure, existing and target customers, service offerings, deployment timeframes, and future technologies. The analysis is an iterative process, which is performed on the customer's premises. A list of preliminary questions should be distributed to the customer prior to the first visit to set the customer's expectations and to assist in defining the design scope.

The documented analysis could be a formal concept of operations for network management or a less formal statement of required features and objectives. However, it should help the network manager to measure success.

This document is the organization's network management strategy and should coordinate the overall business (non-quantitative) goals of network operations, engineering, design, other business units, and the end users.

This focus enables the organization to form the long-range planning activities for network management and operation, which include the budgeting process. It will also provide guidance for the acquisition of tools and the integration path required to accomplish the network management goals, such as SLAs.

This strategic document cannot focus too narrowly on managing specific network problems, but instead on those items important to the overall organization, including budgetary issues. For example:

- Identify a comprehensive plan with achievable goals.
- Identify each business service/application requiring network support.
- Identify those performance-based metrics needed to measure service.
- Plan the collection and distribution of the performance metric data.
- Identify the support needed for network evaluation and user feedback.
- Have documented, detailed, and measurable service-level objectives.

6.5 Document Performance Management SLAs

Service-level management defines and regulates other required performance management processes. Network managers understand that they need capacity planning, but they face budgeting and staffing constraints that prevent a complete solution. Service-level management is a proven methodology that helps with resource issues by defining a deliverable and creating two-way accountability for a service tied to that deliverable. You can accomplish this in two ways:

- Create a service-level agreement between users and the network organization for a service that includes capacity and performance management. The service would include reports and recommendations to maintain service quality. However, the users must be prepared to fund the service and any required upgrades.
- The network organization defines their capacity and performance management service and then attempts funding for that service and upgrades on a case-by-case basis.

In any event, the network organization should start by defining a capacity planning and performance management service that includes what aspects of the service they can currently provide and what is planned in the future. A complete service would include a risk analysis for network changes and application changes, baselining and trending for defined performance variables, exception management for defined capacity and performance variables, and QoS management.

6.6 Document Performance Management OLAs

An operation-level agreement (OLA) often includes hours of operation, responsibilities, authorities, response times, supported systems, etc. OLAs tend to be more technical than SLAs because they define how IT supports IT functions.

Not every SLA requires unique OLAs, and just a few key OLAs can help resolve the silo problem. However, it can be difficult to implement OLAs—especially between departments under different management. Implementing an OLA requires patience and the commitment of all involved, as well as the understanding that each silo has its own job to accomplish. Of course, the common relationship all IT silos share is the provisioning and maintenance of IT services of all kinds to the business.

Just like an SLA, OLAs require monitoring. It is the job of the service-level management (SLM) process to monitor OLAs. If you do not yet have formal SLM, then you must assign an owner to the OLA.

Common OLA contents include:

- Document control and version information
- Authorizations, dates, and signatures
- Objectives and scope
- Definition of the parties of the OLA
- Services covered
- Roles and responsibilities
- Prioritization and escalation
- Response times
- Reporting, reviewing, and auditing
- A list of performance baseline metrics/KPIs

Collect a list of the variables for the baseline, such as polling interval, network management overhead incurred, possible trigger thresholds, whether the variable is used as a trigger for a trap, and trending analysis used against each variable.

6.7 Performance Data Collection

Performance data collection is the process of collecting performance-related management data from the network/system and storing it in a database or data file. The goal of any performance management activity is to collect data that can be used to validate the physical and logical configuration of the network and applications and to localize potential bottlenecks as early as possible.

The polling engine in a network management system can be utilized for data collection purposes. Most network management systems are capable of collecting, storing, and presenting polled data.

6.8 Performance Data Reporting

Performance reporting is the presentation of the collected data. The collected data can be used to analyze faults, growth, and capacity of the network. Reports are useful only if properly correlated data is presented in a useful manner to its intended audience. For example, upper management tends to want reports of the network's performance in simple measurable terms, such as a summary of network availability.

Performance reports can include data on the following:

- Network and device health
- Faults
- Capacity planning

6.9 Measure Network Performance

Performance management is a general term that actually incorporates the configuration and measurement of distinct areas. Four areas of performance measurement in distributed networks are:

- **Measure response time:** Network response time is the time required for traffic to travel between two points. Response times that are slower than normal, seen through a baseline comparison, or that exceed a threshold might indicate congestion or a network fault.
- **Measure accuracy:** Accuracy is the measure of interface traffic that does not result in error and can be expressed in terms of a percentage that compares the success rate to total packet rate over a period of time. For instance, if 2 out of every 100 packets result in error, the error rate would be 2 percent and the accuracy rate would be 98 percent.
- **Measure utilization:** Utilization measures the use of a particular resource over time. The measure is usually expressed in the form of a percentage in which the usage of a resource is compared with its maximum operational capacity.
- **Capacity planning:** As stated earlier, capacity planning is the process in which you determine the likely future network resource requirements to prevent a performance or availability impact on business-critical applications. Refer to the Capacity and Performance Management: Best Practices white paper for more detailed information on this topic.

For instance, the minimum response time for application "x" is 500 ms or less during peak business hours. This defines the information to identify the variable, the way to measure it, and the period of day on which the network management application should focus.

6.10 Perform a Proactive Fault Analysis

Proactive fault analysis is essential to performance management. The data that is collected for performance management can be used for proactive fault analysis. Proactive fault management is the way that the ideal network management system can achieve the goals you determined. The relation to performance management is through the baseline and the data variables that you use. Proactive fault management integrates customized events, an event correlation engine, trouble ticketing, and the statistical analysis of the baseline data in order to tie together fault, performance, and change management in an ideal, effective network management system.

6.11 Performance Baselining

Performance baselining is the process of studying the network, application, and servers/clients; collecting relevant information; storing it; and making the results available for later analysis. A general baseline includes all areas of the network, such as a connectivity diagram, inventory details, device configurations, software versions, application/disk/CPU/memory utilization, link bandwidth, and so on. In summary, the objective of baselining is to create a knowledge base of the network—and keep it up to date.

There are three major components of performance baselining:

- Documenting
- Completing performance assessment
- Understanding and planning

The baselining task should be done on a regular basis, because it can be of great assistance in troubleshooting situations as well as providing supporting analysis for network planning and enhancements. It is also used as the starting point for threshold definitions, which can help identify current network problems and predict future bottlenecks.

Baselining tasks include the following:

- Gather device inventory information. This can be collected via SNMP or directly from the command-line interface (CLI)—for example, show version, show module, show run, show config all, and others.
- Gather statistics at regular intervals.
- Document the physical and logical network, and create network maps.
- Document application- and system-specific details.
- Identify the protocols on your network.
- Identify various applications on your network.
- Monitor statistics over time, and study traffic flows.
- Collect network device/system- and application-specific details.
- Gather database and server- and client-related details.

6.12 Review the Performance Baseline

Network management teams should conduct meetings to periodically review performance reports. This encourages additional feedback, as well as a proactive approach to potential problems in the network. These meetings also provide an opportunity for the planners to receive operational analysis of the baseline and trended data. And operational staff is kept “in the loop” for some of the planning analysis.

Another topic to include in these meetings is the service-level objectives. As objective thresholds are approached, network management personnel can take actions to prevent missing an objective and, in some cases, this information can be used as a partial budgetary justification. The data can show where service-level objectives are breached if proper measures are not taken. Also, because these objectives have been identified by business services and applications, they are easier to justify on a financial basis.

6.13 Document a Risk Analysis Methodology

Perform a network and application risk analysis to determine the outcome of a planned change. Without such an analysis, organizations take significant risks that may hinder success and overall network availability. In many cases, network changes result in congestive collapse causing many hours of production downtime. In addition, a startling amount of application introductions fail and cause impact to other users and applications. These failures continue in many network organizations, yet they are completely preventable with a few tools and some additional planning steps.

You normally need a few new processes to perform a quality risk analysis. The first step is to identify risk levels for all changes and to require a more in-depth risk analysis for higher-risk changes. Risk level can be a required field for all change submissions. Higher-risk-level changes would then require a defined risk analysis of the change. A network risk analysis determines the effect of network changes on network utilization and network control-plane resource issues. An application risk analysis would determine the projected application success, bandwidth requirements, and any network resources issues.

7 Methods to Refine Performance Management

Understanding performance characteristics is fundamental to providing a quality service to the end user as it can help you maintain the performance of network services and applications.

- Detailed performance reporting/information can help you make intelligent decisions that align application, server, and network resources with real business requirements.
- Define performance metrics, SLA, goals, and requirements.
- Define response times, the effects of network bandwidth, network latency and congestion, network errors, node processing, and application behavior.
- Analyze application performance risks and compile results derived from various model scenarios.
- Create recommendations regarding network and application tuning to enhance application performance.
- Review the report and recommendations with the network team to help optimize the performance of applications and network services, and to improve the end-user experience.

8 Broadband Case Study

Service providers that offer xDSL-based broadband services are struggling with the issue of differentiating their services. Offering high-speed access alone, 8 Mbps or greater, is no longer a major differentiator for most customers. Instead offering services that fit the customer's "lifestyle" is much more attractive to today's broadband consumer. In addition newer broadband business models are being developed to allow network owners to revenue-share with over-the-top content providers; but a key requirement will be the ability of the service providers to offer SLAs for their network down to the individual subscriber session.

8.1 Business Requirements

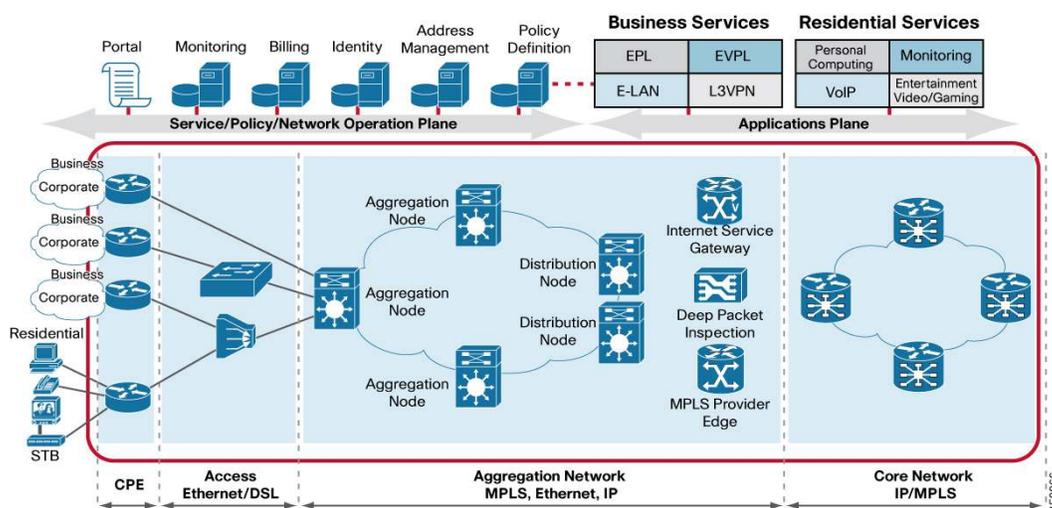
Many broadband service providers use the contention ratio/over-subscription model to define network capacity. This means there is a constant need to monitor the network and understand usage levels. However this approach is not scalable because constant improvements in xDSL technology mean that more and more subscribers can, if they choose, subscribe to higher-bandwidth services. This approach only considers subscribers in terms of aggregation and has no way to cope with individual session requirements. Instead a quality of service (QoS)/quality of experience (QoE)-based approach can be used to do the following:

- Control aggregate usage of specific traffic. For example, control total peer-to-peer traffic in the network.
- Prioritize subscriber traffic over the network based on the commercial service that has been purchased. For example, prioritize gold subscribers' traffic over bronze subscribers' traffic.
- As subscribers become more demanding, broadband service providers will need to offer content-delivery-based SLA's not only to subscribers but also to the over-the-top content providers such as YouTube.

8.2 Typical Architecture of a Fixed Broadband Service Provider Offering

Figure 4 illustrates a typical service provider network supporting fixed broadband services, where performance management can play multiple roles in managing business, residential, and corporate services.

Figure 4. Fixed Broadband Offering Reference Architecture



Key to the overall architecture is the operational plane, which has the following abilities:

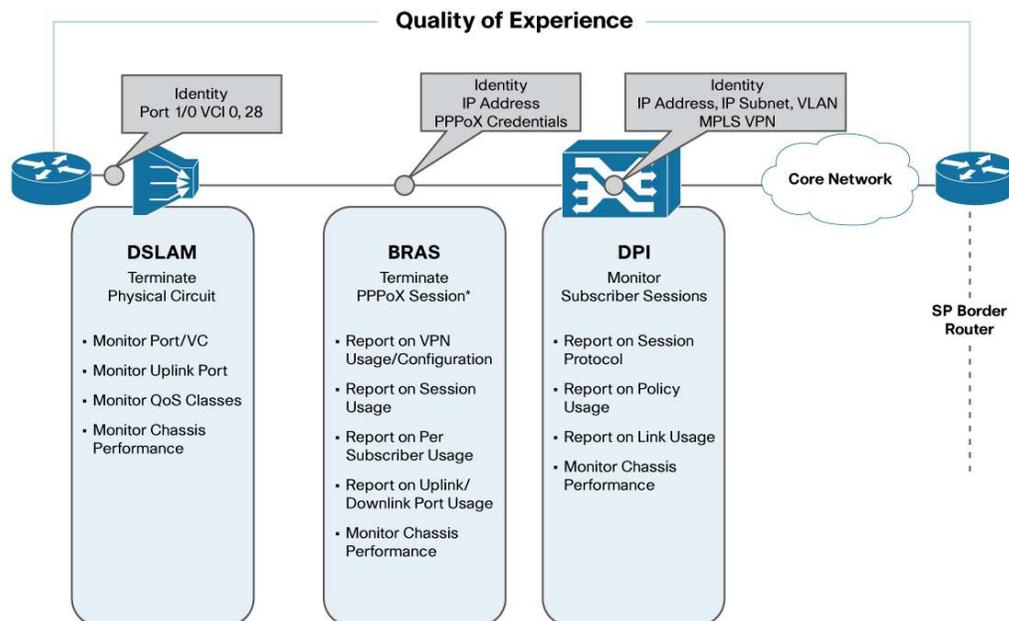
- Enforce individual subscriber policies
- Determine the identity of a subscriber at a given time and provide this information to other systems

A subscriber can be defined as an individual or business that has a commercial relationship with the service provider; in other words, the entity that pays for the service. Associated with the subscriber will be one or many users who actually consume resources on the service provider network. For a residential service, the subscriber and user are typically the same but business services are likely to serve many users and a single subscriber.

8.3 Data Collection Sources and Points

For performance management to work in this context, the performance management systems need to be able to utilize the identity management capabilities being offered by the operational plane to associate data being extracted from network devices. This data typically relates to physical network entities such as a physical port, virtual circuit, IP addresses, etc. to a user and optionally relate that user to a subscriber. The identity of a user on the network is defined as the Network ID. Figure 5 illustrates the main collection points for session data. Users are typically identified at these points.

Figure 5. User Identity and Data Collection Points



* Requirements taken from TR92

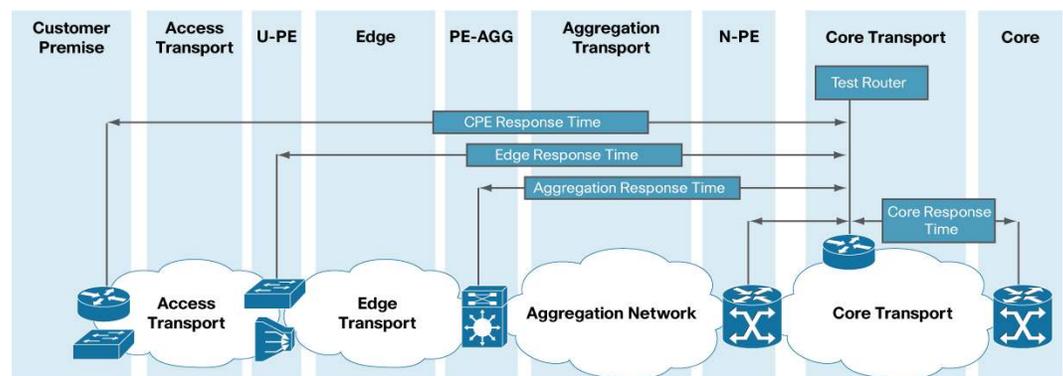
While collecting data at each of the discreet collection points can be a simple process (either using SNMP, RADIUS accounting, NetFlow, etc.), quality of experience is about the end-to-end quality. So the KPIs must be viewed within the context of the user's session. In particular where content SLAs exist, the data must be granular enough to indicate what the user is doing at a given point in time and to provide evidence that the SLA has been met, or indicate if an SLA has been breached. It is also important to understand that "end to end" refers to the path from the customer premises

to the provider border gateway and not the full path to the content because no SLA can be guaranteed over the Internet.

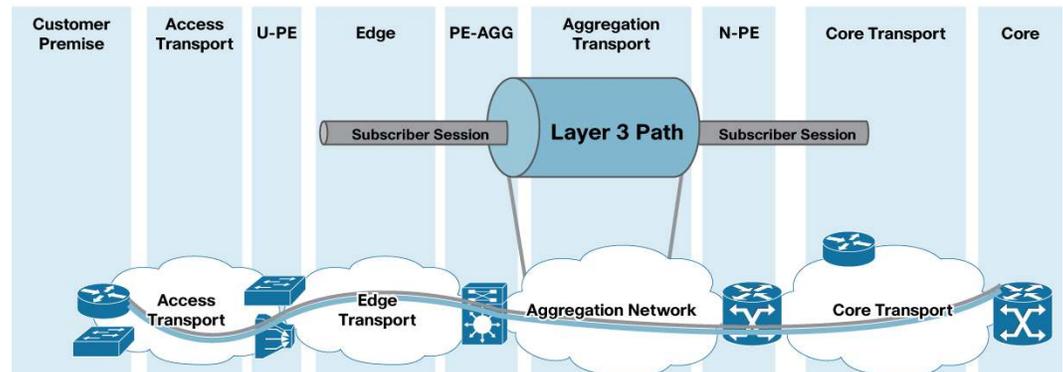
Quality of experience (QoE) can be assessed at either the path or session level. Path-level analysis provides a high-level view of a user's path, hop by hop, from the CPE to the service provider edge. The analysis is fairly simple to implement using standard monitoring technologies but requires each path to be aggregated into a session. Session-level performance allows a much more granular view of the experience an individual user is having but is much more complex to implement and requires Deep Packet Inspection. It is a valid strategy to implement both approaches at the same time. Figure 6 illustrates both approaches.

Figure 6. QoE Assessment

Subscriber Path Performance



Subscriber Session Performance



Extracting the technical measurements is only part of the story. QoE crosses the whole network and so may involve different support groups within the service provider organization. It is important that each of these groups has an operation-level agreement (OLA) that supports the SLA that has been established between the provider and either the end customer or the content provider. This means that each support group must have visibility of the KPIs and associated thresholds, which in turn means that the performance management must be capable of reporting on different aspects of the session and the session path.

QoE is an important part of SLAs between service providers and content providers, and is a service differentiator for residential customers. So SLAs focus much more on the individual subscriber than a service or a network path. QoE can of course only be guaranteed within the service provider's own network. But as content providers such as Google begin to discuss with service providers the business benefits of placing content in situ, then QoE becomes an important consideration. In order to deliver SLAs based on QoE, a new set of KPIs need to be extracted from the network, which may require additional equipment to be deployed such as Deep Packet Inspection devices. Table 2 lists the critical KPIs.

Table 2. Important QoE KPIs

KPI	Granularity	Description
Voice Mean Opinion Score	Subscriber	Provides a numerical measure of the quality of human speech at the destination. This is used as a measurement for subscriber voice QoE.
Video Mean Opinion Score	Subscriber	Provides a numerical measure of the quality of moving picture at the destination. This is used as a measurement for subscriber voice QoE.
Service (P2P, Browsing) Usage	Subscriber	Provides information on the usage of different services such as peer-to-peer (P2P). This is used as a measurement for subscriber usage.
Link Service Distribution	Link	Provides information on the makeup of the traffic passing across the link. This is used to determine what traffic to limit/control to provide a better QoE to the subscribers, such as limiting P2P during peak hours to increase voice/video Mean Opinion Score (MOS).
End-to-End Jitter and Delay	Network	Provides information on the jitter and delay experienced across the network between key points such as POPs to Internet peering points. This is used to ensure that a very coarse-grained level of global QoE can be maintained.
Hosted Service Availability	Service	Provides information on the availability of services that are hosted within the service provider network, including traditional services such as email, caching, DNS, etc. and next-generation services such as video or voice.

These KPIs can be built into OLA's between various internal business units and will form the basis for any SLA's that are delivered to content providers or to residential/business customers. The SLA definitions are typically very enterprise specific and so will not be detailed in this document but best practice suggests that a SLA is more effective when it is described in such a way that the customer understands it and when it is composed of well define SLA components. As well as the traditional SLA components such as availability, integrity etc. performance is normally a key concern to customers and this should be further decomposed into a QoE components as well as standard performance KPI's.

9 Conclusion

In order to deliver high-quality, network-based services, service providers need a strong performance management solution to optimize delivery of key business applications and to help their customers manage IT expenditures and achieve a rapid return on investment.

Traditionally performance management systems have been internally focused, providing views of network health. If service providers wish to participate in the web 2.0 value chains, these systems need to go deeper into the traffic passing across the network and look at an individual user's session and determine the quality of experience a user is having. This will allow the service provider to differentiate their services with both end users and content providers. Because this is a technically challenging requirement for most service providers, the systematic approach described in this paper should be used to simplify the challenge.

This is particularly critical in Broadband service providers as this market is fast reaching saturation point and service providers need to differentiate their service offerings from the competition by focusing on providing good Quality of Experience thus reducing churn and attracting new subscribers. Quality of Experience (QoE) needs to be provided at a more granular protocol level than TCP/IP, it needs to be provided at the Peer to Peer (P2P), Instant Messaging level as so a key requirement for service providers wishing to deliver QoE metrics is the ability to identify their subscribers on the 'wire' and classify this traffic down to required level of granularity.

10 References

- [Performance Management: Best Practices White Paper](#)
- [Performance Reporting Concepts and Definitions—TMF701 v2.0](#)
- [An Extension to the OSI Model of Network Management for Large-Scale Collaborative Performance Measurement](#)
- [Cisco IOS IP Service Level Agreement Data Sheet](#)
- [Cisco Service Control](#)



Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV
Amsterdam, The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

CCDE, CCENT, Cisco Eos, Cisco Lumin, Cisco Nexus, Cisco StadiumVision, the Cisco logo, DCE, and Welcome to the Human Network are trademarks.; Changing the Way We Work, Live, Play, and Learn is a service mark; and Access Registrar, Aironet, AsyncOS, Bringing the Meeting To You, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, CCVP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Collaboration Without Limitation, EtherFast, EtherSwitch, Event Center, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, iQ Expertise, the iQ logo, iQ Net Readiness Scorecard, iQuick Study, IronPort, the IronPort logo, LightStream, Linksys, MediaTone, MeetingPlace, MGX, Networkers, Networking Academy, Network Registrar, PCNow, PIX, PowerPanels, ProConnect, ScriptShare, SenderBase, SMARtNet, Spectrum Expert, StackWise, The Fastest Way to Increase Your Internet Quotient, TransPath, WebEx, and the WebEx logo are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or Website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0805R)