Service-Level Management

Definition

Service-level agreements (SLAs) are contracts between service providers and customers that define the services provided, the metrics associated with these services, acceptable and unacceptable service levels, liabilities on the part of the service provider and the customer, and actions to be taken in specific circumstances. Service-level management is the set of people and systems that allows the organization to ensure that SLAs are being met and that the necessary resources are being provided efficiently.

Overview

This tutorial is a guide for companies looking to manage not just network equipment but the services associated with that equipment. The tutorial introduces service-level management, its relationship to SLAs, and the requirements of a service-level management system.

Topics

Introduction
Background
What Is Service-Level Management?
What Is Needed?
The Solution
Self-Test
Correct Answers
Glossary

1. Introduction

On May 18, 1998, an estimated 80 to 90 percent of the pagers in the United States went out of service when the \$250-million Galaxy IV satellite rolled out of control and lost track of the Earth. PanAmSat, the Greenwich, Connecticut

company that operates Galaxy IV, switched communications to another satellite, but not before pager, radio, television, and credit-card transmissions were disrupted for up to two days in some areas of the country. The net economic cost was astronomical.

When America Online (AOL) went to a flat-fee, unlimited-access rate, this Internet service provider (ISP) lost thousands of subscribers and offered refunds to countless others. The change had increased network traffic so much that callers to AOL had difficulty making connections, especially during peak periods. The network congestion also made it harder for non–AOL callers to make connections, because they were using some of the same telecommunications infrastructure.

These real-world examples point to the growing challenges involved in managing service availability and service levels in today's explosive and rapidly changing communications environment. They represent worst-case scenarios not only for end-user customers, but also for those organizations responsible for managing service levels and ensuring quality of service (QoS).

2. Background

Managing communications service delivery is challenging. Typical environments consist of equipment from several vendors, supporting multiple protocols, with multiple services and multiple users. The network and equipment performance is dynamic, the amount of data is large, but the amount of useful information is limited. In this complex environment, it is difficult to detect potential problems before they surface. Once they do surface, it takes time to discover the true root cause, determine the best solution, measure the potential business impact, and take corrective action. The more time it takes, the greater the negative impact on service availability and overall service quality.

While managing service delivery is challenging, the cost of not doing it well is enormous. In today's world of converging voice and Internet services, which are offered over a combination of wired, wireless, and cable media, the service providers that cannot keep up will be left behind by those that can differentiate their services and get them to market quickly.

The task of building network operations centers (NOCs) and operations support systems (OSSs) to manage effectively the highly complex, distributed, missioncritical networks used to deliver today's data and voice services is one of the most important issues facing service and equipment providers. Organizations must integrate legacy systems with emerging technologies, manage a host of diverse and often incongruous technologies, and deploy flexible, platform-independent applications—all with minimal staff. And they must do so faster than their competitors to ensure the successful launch of new products and services. Whether applying Internet or corporate Intranet, data transmission, fax, wireless, landline telephone or cable service, users expect virtually 100-percent availability. To deliver it, communications service providers must be able to anticipate potential network problems before they occur. When interruptions do happen, they must be able to isolate the cause immediately and take appropriate corrective action. Failure to meet these objectives will cost companies existing and potential customer revenue.

This increasingly demanding environment has created OSS requirements that far exceed the capabilities of first-generation, off-the-shelf software applications. These applications are not designed to integrate the broad spectrum of new technologies with existing legacy systems, manage the multitude of continuously changing variables in real time, or scale to the ever-increasing traffic volumes generated today.

A typical service provider NOC receives hundreds of events per second from tens of thousands of managed resources. Typical OSS software includes faultmanagement event browser tools, with exception messages presented to the operators in a standard form and grouped by operator responsibilities (management domains and functional areas). Those messages may be filtered, and simple automations, such as event stream deduplication or managed resource up-down correlation, may be performed.

However, when complex problems occur (as is the norm), operators and knowledge experts are faced with daunting problem-solving tasks such as the following:

- correlating multiple events from one or more managed resources across one or more management domains
- running diagnostic tests
- calling on the knowledge experts
- isolating the problem and its root cause(s)
- eliminating symptomatic errors
- determining the impact on users and services
- notifying appropriate management and customer support centers
- applying corrective actions to resolve each problem
- taking the failing component(s) out of production for repair
- submitting trouble tickets

The situation is typically compounded by adding new managed resources to the production environment at a rate of several hundred per week. Coordinating the change/fault management responsibilities and actions is a challenge. An automated method of incorporating new managed resources and removing retired managed resources from the change/fault/problem/asset management-system environment is essential.

Also, managed resources are periodically removed from production for routine maintenance, upgrades, and problem fixes. The exception reporting, problem resolution, and notification procedures vary according to the state of the managed resource.

Isolating and resolving these problems is often a manual, labor-intensive and time-consuming process—even if it is in the order of minutes. Impacting service availability even for seconds can be significant in terms of customer satisfaction and lost revenue.

3. What Is Service-Level Management?

Once again, SLAs are contracts between service providers and customers that define the services provided, the metrics associated with these services, acceptable and unacceptable service levels, liabilities on the part of the service provider and the customer, and actions to be taken in specific circumstances. The service provider may be a traditional carrier or another group in the company that provides the network services. Service-level management is the set of people and systems that allow the organization to ensure that SLAs are being met and that the necessary resources are being provided efficiently. In the telecommunications management network (TMN) model, service management logically resides above the element and network management layers and below business management (see *Figure 1*). SLAs are critical success factors in terms of differentiating and ensuring service delivery.



Figure 1. Service Management's Role in the TMN Model

Web ProForum Tutorials http://www.iec.org The definition of the SLA is important in terms of defining customer expectations. Most SLAs are developed after first baselining existing services and then by defining what, if any, new services are required. An SLA is intended to be an objective tool that helps service providers define benefits versus cost tradeoffs and deliver communication services that provide the best value to the customer.

Traditional approaches to service management involving event filtering, correlation, and even root-cause analysis, simply fall short. They cannot handle the dynamic complexity of today's networks and cannot automate the labor-intensive, time-consuming steps required to investigate and correct each potential problem. Some products do a good job of displaying data, but most of the burden still falls on the operators, who must still perform all tests, diagnose the cause, measure the business impact, and then decide on and execute an appropriate response. The biggest challenge is to ensure that one measures and manages the parameters that matter (not just those that are readily available), meaning those that impact the service as defined in the SLA. This means being able to collect low-level equipment, network, and application data and combine it into service data, given that single, easily measured and managed service attributes do not often present themselves.

4. What Is Needed?

To manage service delivery well, it is essential that the service management system work well with the network and element management systems to handle each of the following tasks:

- **dynamic network representation**—to handle new and out-of-service managed resources
- **proactive data and event monitoring**—to integrate data from a set of diverse vendor and protocol sources and identify potential problems
- **automated topological and model-based reasoning**—to understand the impacts of events based on connectivity and configuration
- **cause-effect inferencing**—to map between the effects of anomalies and their service impacts
- **root cause analysis**—to differentiate between symptoms of a problem and the true problem
- **operator guidance**—to assist to prioritizing operator actions

- **automated testing**—to verify diagnoses while minimizing network impacts
- **automated fault correction**—to implement network reconfiguration to optimize service delivery

These capabilities are essential for truly effective service-level management. Without them, most of the event management responsibility continues to fall onto already overburdened operators.

Mapping network events onto services and related SLAs allows for real-time management of the business. Key performance indicators can be polled periodically, and trend analysis of the historical values can be used for proactive SLA management, predicting limit violations. As a consequence of a violation of an SLA condition, any actions—not just reports—can be performed, allowing corrective measures to be automated. The service management system should include built-in object models of the network devices and their connection, the services, SLAs, and customers. It should be able to perform what-if analysis for possible customers affected by a potential failure and be able to generate a list of all devices used by a customer or group (see *Figure 2*).

Figure 2. Basic Network, Service, and Customer Mapping



5. The Solution

The actions performed by a well-designed service management system include the following:

- understanding the complex relationships between each object, process, and event
- anticipating and diagnosing events based on this understanding
- automating all necessary tests and data gathering (such as transaction language 1 [TL1] messages, signaling network management protocol [SNMP] gets, and Internet protocol [IP] echo pings)

- automating corrective actions
- measuring the potential business impact of each event (on service-level agreements, for example)
- accepting data from any source
- providing a flexible user interface (including Web and Java-based interfaces)

This functionality in an OSS is achieved by integrating software with capabilities at three levels (see *Figure 3*):

- support for the required managed object interfaces (event integration)
- powerful reasoning on this collected event data to identify the root cause of problems and service impacts (root-cause analysis and service impact assessment)
- support for complementary best-of-breed applications (application integration)



Figure 3. OSS Event Integration, Processing, and Outputs

The service management system must accept inputs from a diverse set of managed objects. These include raw feeds from ASCII ports, control interfaces to switches, standard management interfaces with SNMP or common management information protocol (CMIP)—manageable devices, as well as newer common object request broker architecture (CORBA)—based interfaces. All of this information must be combined intelligently, in real time, to provide the diagnosis of current service-level impacts, predicted service-level impacts, and necessary actions to ensure that service-level expectations are met. Finally, integration with other applications, such as the billing system, is required to reflect the impacts of the services delivered accurately.

The Result

Successful OSS implementers today must turn their focus from simply managing the equipment and networks in their businesses to managing the services supported by those networks and equipment. The fact that a given piece of equipment is failing may or may not be important. If that piece of equipment is supporting a current videoconference for the CEO of a customer, for example, it is quite important. A well-designed service-level management system will provide the rapid diagnosis and response to each potential problem that supports improved QoS and resource allocation, escalation, and priorities.

Self-Test

- 1. Which of the following is not characteristic of communications service delivery environments?
 - a. dynamic network performance
 - b. dynamic equipment performance
 - c. limitless useful information
 - d. large amounts of data
- 2. Which of the following is not true of organizations that seek to manage their networks more effectively and improve service delivery in today's environment?
 - a. They employ an extensive staff.
 - b. They deploy platform-independent applications.
 - c. They manage diverse, often incongruous technologies.
 - d. They integrate legacy systems and emerging technologies.
- 3. First-generation, off-the-shelf software applications were designed to scale to ever-increasing traffic volumes.
 - a. true
 - b. false

- 4. Exception reporting, problem resolution, and notification procedures vary according to the state of the managed resource.
 - a. true
 - b. false
- 5. A few seconds of lost service availability will not affect customer satisfaction and revenue.
 - a. true
 - b. false
- 6. In the TMN model, service management logically resides _____ the element and network management layers and _____ business management.
 - a. above: below
 - b. below: above
- 7. In today's typical service management system, most of the burden for turning operational data into information falls on _____
 - a. products
 - b. traditional service management approaches
 - c. operators
 - d. network managers
- 8. The service management system should be able to generate a list of all devices used by a customer or group.
 - a. true
 - b. false
- 9. Which of the following does not accurately describe a well-designed service management system?
 - a. understands complex relationships between each object, process, and event
 - b. automates corrective actions

- c. provides a flexible user interface
- d. accepts data from a limited number of sources
- 10. The functionality in an OSS is achieved by integrating software with capabilities at ______ level(s).
 - a. one
 - b. two
 - c. three
 - d. four

Correct Answers

- 1. Which of the following is not characteristic of communications service delivery environments?
 - a. dynamic network performance
 - b. dynamic equipment performance

c. limitless useful information

d. large amounts of data

See Topic 2.

2. Which of the following is not true of organizations that seek to manage their networks more effectively and improve service delivery in today's environment?

a. They employ an extensive staff.

- b. They deploy platform-independent applications.
- c. They manage diverse, often incongruous technologies.
- d. They integrate legacy systems and emerging technologies.

See Topic 2.

3. First-generation, off-the-shelf software applications were designed to scale to ever-increasing traffic volumes.

a. true

b. false

See Topic 2.

4. Exception reporting, problem resolution, and notification procedures vary according to the state of the managed resource.

a. true

b. false

See Topic 3.

5. A few seconds of lost service availability will not affect customer satisfaction and revenue.

a. true

b. false

See Topic 3.

6. In the TMN model, service management logically resides ______ the element and network management layers and ______ business management.

a. above; below

b. below; above

See Topic 3.

- 7. In today's typical service management system, most of the burden for turning operational data into information falls on ______.
 - a. products
 - b. traditional service management approaches
 - c. operators
 - d. network managers

See Topic 3.

8. The service management system should be able to generate a list of all devices used by a customer or group.

a. true

b. false

See Topic 4.

- 9. Which of the following does not accurately describe a well-designed service management system?
 - a. understands complex relationships between each object, process, and event
 - b. automates corrective actions
 - c. provides a flexible user interface

d. accepts data from a limited number of sources

See Topic 5.

- 10. The functionality in an OSS is achieved by integrating software with capabilities at ______ level(s).
 - a. one
 - b. two
 - c. three
 - d. four

See Topic 5.

Glossary

AOL

America Online

ASCII

American Standard Code for Information Interchange

ICMP Internet control message protocol

IP Internet protocol

ISP Internet service provider

NOC network operations center

OSS operations support system

ping packet Internet groper

QoS quality of service

RAM random access memory

SLA service-level agreement

SNMP signaling network management protocol

TL1 transaction language 1

TMN telecommunications management network

WAN wide-area network