Achieving customer-centric assurance

An HP white paper on NGOSS

Table of contents

Abstract: With the pressure to gain cost advantage and service leadership, communications service providers (CSPs) can no longer rely on a network-centric, or in the best case, service-centric approach to conduct their business. An approach where the perception and experience of customers define the business strategy is required, taking into account all the facets of the interaction between the service provider and the customer.

In HP Communications & Media Solutions (CMS), we are leveraging our extensive experience in telecom and IT to help CSPs transform the customer experience to best meet the business challenges of today and tomorrow.

Too often today, customer experience management (CEM) solutions are only focusing on detecting customer experience problems and not addressing the areas of identification and resolution. In HP Next-generation Operating Support System (NGOSS) Solutions, we offer innovative customer-centric assurance capabilities to address the needs of the full CEM process covering: detecting quality of experience problems,[1] identifying them and their root cause, facilitating their resolution, and enabling ongoing preventative measures.

Using a methodology that drives the technology from the business requirements, this paper will expose the concerns leading to the business case for a customer-centric operations support system (OSS), develop what is needed to address these concerns, and finally take a closer look at how the solution should look to ensure it meets the business goals.

Executive summary

Communications service providers (CSPs) are seeking ways to gain cost advantage while maintaining service leadership to attract new customers and retain existing ones.

In the past, assuring the availability and quality of the network proved adequate to ensure high quality of the services delivered to end customers. Today, with the large number and variety of different services accessed in multiple ways by customers using different devices, this approach is no longer sufficient. Moving to a service-centric way of operating is a step in the right direction, but still not addressing the fundamental business problem of why customers change or stay with a service provider, and what attracts new customers. A different approach is required, where the actual perception and experience of customers define the business strategy of the service provider.

The result is the transformation we see today in CSPs, involving both network infrastructure as well as organizational changes, in order to focus on the customer, not just the network or services.

The need to manage the customer experience end to end puts new requirements on how the operations support system (OSS) and business support system (BSS) should support the business strategy and the service provider's business processes, including provisioning, assurance, and billing. Customer experience management (CEM) addresses the need for taking into account all the facets of the interaction between the service provider and the customer.

Customer experience management is a broad concept—in fact, as well as more easily measurable factors such as activation time for services, service quality, speed of repair, and length of hold time for customer service, it also encompasses an emotional dimension. Factors such as the perceived technical level of customer care, badly performing voice-recognition systems, and design and ease of use of Web sites all enter into the appreciation customers will have of people, services, or products representing the service provider organization.

In this paper we focus on one essential part of CEM: measuring and providing that each customer receives the expected or contracted quality of the service he/she has subscribed to, i.e., the OSS aspects of CEM and, in particular, how the service assurance business process needs to be adapted to support customer-centric operations.

Solution considerations

The customer experience management process

As we traverse the process used to address a typical customer experience problem,[1] we encounter a spectrum of needs:

- How do we detect a potential customer experience problem?
- What is needed to identify the problem?
- How do we fix the problem?
- What can be done to prevent the problem from happening again?

1. We use the term "problem" with the traditional telecom meaning. Readers familiar with Information Technology Infrastructure Library (ITIL) terminology should substitute "incident."

In the context of customer-centric assurance, multiple techniques and tools are used to respond to these needs, focusing on the OSS specifically:

- Problem detection
  - Active probes (tools simulating user activity)
  - Passive probes (such as measuring tools residing in customer device, on network signaling links, or specific passive "listening" tools such as deep packet inspection)
  - Usage data
  - Network and service fault and performance management systems
  - Service quality and SLA management systems
- Problem identification
  - Problem correlation and root-cause analysis
  - Service quality and impact analysis
  - Usage and probe data pattern detection
- Problem resolution
  - Problem diagnosis and testing
  - Workflow automation/dispatching—trouble ticketing
- Problem prevention[2]
  - Network and service thresholds and trend analysis

These points will form the basis for discussing the concerns and requirements leading to a customer-centric assurance solution.

2. Using ITIL terminology, "problem management."

Figure 1. Customer experience management process

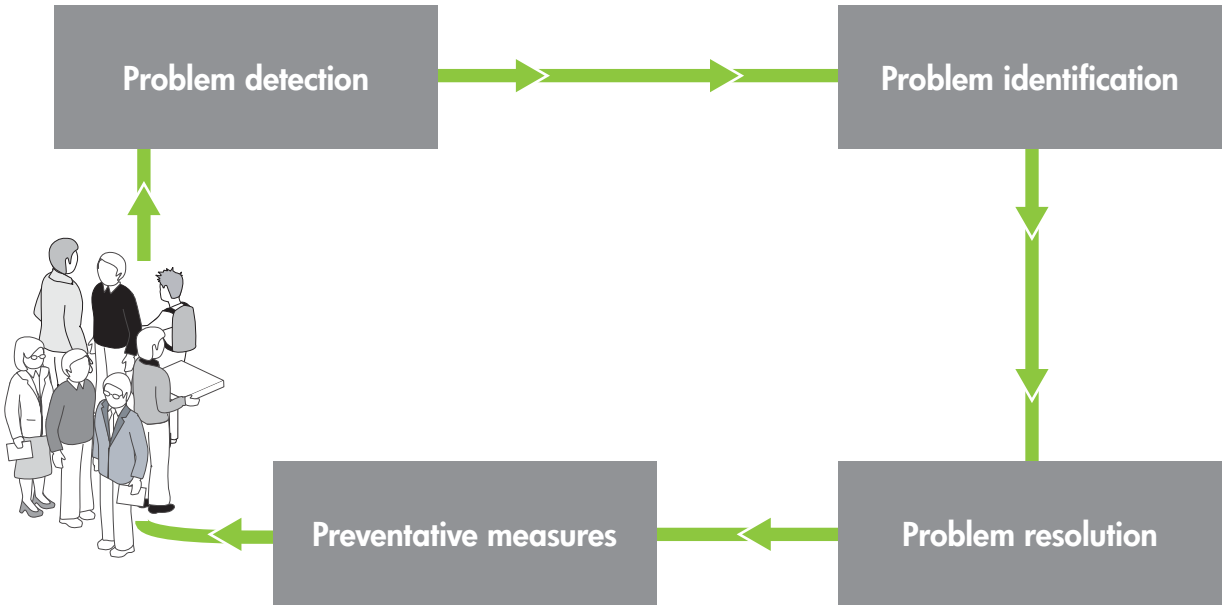

Concerns and challenges in CSPs today

A customer-centric assurance solution

As mentioned above, an ideal customer-centric assurance solution should address a full CEM process covering: detecting quality of experience problems, identifying them and their root cause, facilitating their resolution, and ensuring ongoing preventative measures. Too often today, CEM is only focusing on quality of experience (QoE) problems and not addressing the areas of identification and resolution.

We have looked closely at current processes and tools in OSS and identified three key areas for improvement:

1. **Problem correlation and root-cause analysis (RCA)**

   Quickly identifying a network resource problem and understanding its impact on services is the first step toward ensuring customer satisfaction. Unfortunately existing approaches often fail to keep up with change, are difficult to maintain, and fail to address cross-technology problems. A radically different approach is required where the network structure is captured in a dynamic topology model that accurately reproduces the behavior of the real network.

2. **Integration across resource and service layers**

   Silo organizations, process, and applications prevent the move to customer-centric operations. The lack of consolidation across multiple network domains and equipment vendors may successfully be addressed by a manager of managers approach. But more often, there is a lack of integration across vertical layers. Consolidating the resource and service management layers allows for linking network with service operations and enables efficient problem detection, identification, and resolution, resulting in increased overall operations efficiency.

3. **End-to-end service quality management**

   A comprehensive customer-centric assurance solution must measure the customer experience by collecting metrics through passive probes, active probes, and usage data. But being able to detect a customer experience problem is not sufficient, the solution must also provide the capability of correlating these metrics with other data sources, such as transaction and session statistics, network and services data, along with business-related data, in order to also identify and resolve the customer problem.

The next sections will develop the concerns and challenges related to these three areas, and what is needed to address them.

Attaining cost-effective problem correlation and root-cause analysis (RCA)

Service can be defined as how businesses are measured by their customers, be it the choice of service offered, its quality, its reliability, or cost. It is the most visible and obvious means that a customer has to gauge the performance of an organization as a whole. Any degradation of the quality of service is potentially noticeable by end customers and could lead to a negative impact on the business through loss of revenue, poor customer relations, financial penalties, and increased customer churn.

In the telecommunications domain, a variety of technologies and infrastructure are used to deliver the service. As communication networks are frequently changing in response to market forces and the introduction of new technologies, their complexity and the interaction between devices and services make it difficult to identify the actual cause of a problem and its impact on services.

When a failure occurs, a significant amount of data may be generated by the network to inform the operations staff about the problem. This is because alarms get generated not just from the point of failure, but also from affected network equipment and services downstream within the network, resulting in so-called sympathetic alarms. As a result, a huge number of alarms can be delivered to the network monitoring team(s) within a very short period of time.

Identifying the root cause of a problem is therefore typically a time-consuming process, as the network operators have to work through a large number of alarms to try to identify the problem's origin. Once identified, operators then need to put in place a process to ensure its resolution and determine affected customers. In an attempt to improve the situation, CSPs have used a number of different approaches coming on the market in the last ten years.

Examples are rule-based, case-based, and several other correlation technologies for automatically detecting and handling underlying network problems. These systems typically maintain a knowledge base of failure scenarios and process large quantities of event information from the network, in an attempt to detect a pre-defined failure pattern.

Rule-based systems are typically well suited for specific technology domains, and less adapted to addressing the interactions between layers and network domains, or network environments that often change. Ideally, rule-based systems allow for automating well-defined, frequently performed tasks in operations.

Unfortunately, some of the other approaches such as case-based systems have a number of drawbacks when trying to address rapidly changing environments:

- The solution is inflexible. When the topology or inventory is changed or a new service is implemented, the system has to be re-engineered by skilled personnel. New interactions between services and devices may be introduced that will need to be well understood to ensure accurate, reliable information is provided. This is a time-consuming, expensive, and error-prone process that requires extensive testing prior to deployment.
- There is a requirement for specialist skills to configure the associated cases, as the information has to be correctly entered in a form that accurately translates the signature of a failure into a syntax that the system can understand.
- This complexity is difficult to manage effectively on a day-to-day basis.
- These systems are typically not suited or capable of handling cross-domain correlation scenarios, i.e. those involving several technology domains, such as radio and fixed-line access networks, IP core network, and optical transmission networks.

- Due to the complexity and cost of maintaining the system, the view of the network built into the original solution slowly—but surely—diverts from the actual network topology, eventually leading to operator staff abandoning the system.

In order to overcome these drawbacks, a radically different approach is required. A solution is needed where the propagation of failures in the network and connectivity information already exist inside the system in a pre-built, automatically maintained, dynamic, topology model that is easily modifiable and truthfully mimics the behavior of the real network.

The lack of cross-layer consolidated operations

CSPs have understood that silo organizations, process, and applications prevent their move to customer-centric operations. For example, multiple, non-integrated applications from multiple vendors residing in the element, resource, and service layers, inhibit efficient operations. Organizational and process consulting can aid with best practices for transformation along these lines.

In the application area, the traditional approach to consolidate multiple network management silos consists of adding a horizontal layer above the different domain managers and hence creating the concept of manager of managers. This approach has proven successful in numerous CSP projects, resulting in increased operations efficiency and significant cost savings.

An area that has been less explored is that of vertical consolidation, i.e. a system capable of integration across the element, resource, and service management layers. Different object and topology views in non-integrated databases, and a disconnected resource/service view make the problem resolution process a challenge for operations personnel.

Many service providers today are organized with separate network and service operations staff. To avoid manual correlation between teams, an ideal approach should consolidate the horizontal service and resource layers in an integrated end-to-end view of all important network links and elements, as well as systems and applications, and at the same time the state of dependent services. This enables operators to reduce the time spent in diagnosing and resolving cross-layer, resource-service-customer related problems.

When a problem in the network resource layer has an impact on a customer service, the monitoring system must promptly notify the operations staff, which will urgently need to identify the problem in order to initiate repair actions. It's essential that the problem be identified before the customers notice it. In the worst case, customer care should at least be able to tell the customer who calls, "We are aware of the problem and are working on it," and ideally, "it will be fixed within x hours." Failure in doing so highly increases the risk of customer churn—and it is less costly to retain an existing customer than to acquire a new one. Ideally, the tool should allow users to start from the high-level customer impact notification and be able to drill down into any lower resource level to precisely investigate the set of alarms ultimately responsible for the service object's failing state.

But the end-to-end view is just the visible part of the application. The integration between objects should happen through a shared service object repository, capable of supporting requirements for convergent telecom and IT services, while adhering to the Shared Information Data (SID) [GB922] model from the Telecommunication Management Forum (TMF).

Achieving true end-to-end service management

There are real gains to be made for service providers equipped to profit from emerging opportunities. Success is going to those able to roll out innovative new services quickly and with consistently high quality. The challenge is significant, but so is the payoff—better customer adoption of new services, improved top-line and bottom-line revenue, greater insight and control over services portfolio management, and powerfully enhanced customer satisfaction.

Service quality is a deciding factor in determining customer satisfaction. CSPs need to resolve service quality problems before they negatively affect the customer experience. This requires transparency into the services value chain from end to end. Although essential, monitoring the network for fault and performance problems is not necessarily indicative of the performance perceived by the customer.

Since the service experience end users have is of paramount importance in the competitive marketplace, a comprehensive customer-centric assurance solution must provide an end-user experience view allowing service operators to get an end-user perspective on service levels, which would help realize a tight alignment between service operations and business goals. Services and customer experience views must be monitored in real time for adherence to defined and measured key quality indicators throughout the service lifecycles. Measuring the customer experience involves collecting metrics through passive probes, active probes, and usage data.

However, this is not sufficient. The solution must also provide the capability of correlating multiple data sources, including customer experience metrics, transaction and session statistics, network and service fault and performance data, with business related data, such as call center and order handling performance, in order to not only detect, but also be able to identify and resolve, a problem.

Existing OSS approaches to managing customer experience often only address some of these aspects and thus fail to offer a comprehensive end-to-end solution.

A few words on OSS integration and unified data management

To be able to share data across applications, we need consistent information models. Common syntax and semantics are crucial for communication to take place. By defining a consistent model based on specifications such as TMF's SID, mentioned earlier, information models can be easily extended and any application that needs to use the information can easily understand it.

Unified data management is one of the most critical components of a successful NGOSS. In the IT Infrastructure Library Version 3 (ITILv3), the IT industry has come to realize something that OSS theory has known for a while. A federated repository of the management information is the only way to provide the necessary balance among completeness of information, access to information, and accuracy of information, while allowing the focus to move from technology to business outcomes. This move by ITIL comes at an opportune time, allowing the network focus of traditional telecom inventory systems to be enhanced with the application and services focus of IT configuration management database (CMDB) systems. By federating these together, it becomes possible to have a comprehensive view of the end-to-end service, augmenting the network inventory with the complex relationships introduced by value-added services.

As a result, the new federated repository would be a combination of:

- Traditional inventory, capturing information about all the network and IT resources, and how they are connected
- Services inventory, which maintains information on how the individual service instances are provisioned

The solution—how should it work?

Coming back to process used to address a typical customer experience problem, let's see how the proposed key OSS improvements, together with existing processes, will be exploited in a typical use case.

Detecting the problem

The event that triggers the process at the very beginning would preferably be a notification to the network or service operations staff that something is not working as expected, before any customer notices the problem. For example, the notification could be an e-mail, an SMS, a new trouble ticket, or in the form of alarms displayed in a fault management application window. In the worst case, the operator personnel would not become aware of the problem until a customer complaint is logged by customer care, typically resulting in a customer service trouble ticket that will be used to follow the customer problem resolution process from opening to closure.

With the assumption that the CSP is using state-of-the-art customer-centric assurance processes, the service operations staff will be the first to get notified of a potential degradation of the quality of experience (QoE). The service quality management tool uses visual indicators like icons changing colors to indicate a deviation or degradation in service levels, monitored in real time and compared against appropriate thresholds. At the same time, notifications, such as those listed above, are sent to alert the staff to take action before services level agreements are violated or the customer experience suffers.

Optionally, the notification could carry a priority according to the most profitable services and customers, derived from comparing collected metrics against individual customer service level objectives, which would allow the operations team to focus on the most critical issues.

We assume the customer experience metrics are collected using a combination of active probes, simulating and measuring the user activity, and passive probes listening to the signaling and traffic channels, and analyzing protocol packets. The metrics, in the form of key performance indicators (KPIs) are then used to compute key quality indicators (KQIs), which in turn are validated against service level objectives (SLOs) defined in operational or customer SLAs.

At approximately the same time, the network operations staff is receiving other notifications, due to a large number of alarms arriving through fault and performance management systems. In the same way as the customer experience metrics, selected alarms (those that have a potential service impact) get mapped to KPIs, and enter into the computation of KQIs used for validation against service level objectives.

The large alarm volume is because alarms are not just generated from the point of failure (symptomatic alarms) but also from affected devices and services (sympathetic alarms). Further, performance degradations are likely to be noticed by other OSS systems, e.g. performance management systems, which will generate additional threshold-crossed alarms.

As a result, we are facing at this point a large volume of outstanding alarms seen by the network operations staff, while at the same time multiple indications of degraded QoE in the service operations center.

Identifying the problem

With traditional systems, a laborious process involving different operations teams would now commence to try to figure out which part of the network, and eventually which component(s) are responsible for the quality degradation seen from collected customer experience metrics.

In our case, network operations will use a two-step approach: first, finding the correspondence between the customer service degradation and underlying network resources, then performing an automatic root-cause analysis from the service impact alarms associated with the network resources to identify the origin of the problem.

To rapidly recognize service problems and initiate repair actions, the application needs to provide real-time monitoring of service availability. It should provide an intuitive view of the services, and how they relate to each other, via a graphical service tree, continuously computing and propagating service component status using the data collected from the resource layer. Additionally, to be able to react before any noticeable customer impact, the compliance with service levels needs to be monitored in real time, using thresholds set in such a way that a degradation of the service quality can be proactively recognized.

In order to efficiently relate resource problems to affected services, the same graphical view also provides a view of the underlying resources with their outstanding alarms. The link between resource alarms and service statuses would be done by qualifying certain resource alarms as service impact alarms, map these into key performance indicators (KPIs), and associate them to the related service objects. Their statuses would then be propagated against the service tree.

Network operators will be looking at a subset of the service views seen by the service operations staff, with an additional network resource view, which allows them to identify a dependency between the partial degradation of the concerned customer service and underlying network resources. This is done by drilling down from the degraded service component in the service tree view to the corresponding network resource component. In fact, within the large volume of outstanding alarms, just one is marked as quality of service alarm, associated with the network resource, and indicating performance degradation. It was generated by the problem correlation and root-cause analysis system to indicate service impact at the network resource level and will now allow the operator to quickly display the corresponding origin of the problem—the root-cause alarm.

Any topology-based correlation system needs to work in concert with existing alarm and performance systems as an intelligent information condenser—taking in high-volume, low-value information from source systems and sending high-value, low-volume notifications to operators and other systems. Our quality of service alarm is an example of such a value-added alarm. Needless to say, any such system should also be architected for scalability and distribution and ideally use data-driven techniques throughout to minimize overall system integration costs.

The topology-based approach allowed us to correlate alarms across multiple networks and technologies to determine service impact—used to feed the state propagation in the service tree—as well as automatically perform root-cause analysis to identify the origin of the problem.

Fixing the problem

Continuing our use case, a workflow to fix the problem was also automatically initiated. A network trouble ticket was automatically opened to handle the root-cause alarm and dispatched to the relevant network operations team through a simple alarm action rule. The person on duty in network operations verifies the contents of the ticket and takes responsibility of the ticket, since the problem is related to his/her domain. The state of the problem and the person in charge of the ticket are visible across the operations teams.

At the same time, an associated customer service ticket is also created, in order to allow customer care to answer to any customer calls in case there is a noticeable impact on customer experience before the problem gets repaired.

Sometime later, once the operator repair team has fixed the cause of the problem, several events occur that will bring us back to normal. New, collected probe data will result in clearance of the degraded state of the customer service, clear alarms are sent from related network resources and elements, the problem correlation system will re-evaluate service impact, and finally, states of alarms and trouble tickets are synchronized, resulting in manual, or optionally automatic, closure of the network and service trouble tickets.

In order to enable the different OSS components to interact seamlessly in the way outlined above, data needs to be shared across applications through consistent information models. For example, by defining the model based on specifications such as TMF's SID, we enable applications that need to use the information to easily understand it, while allowing convergence of telecom and IT through the use of configuration management database (CMDB) systems, as described in the previous chapter.

Preventing the problem from happening again

To be able to enhance customer perception in the longer term, and drive new service offers, as well as feeding important service performance information to higher-level management, the customer-centric assurance solution needs to continuously store all real-time metrics, KPIs, and service-level incidents (degradations and violations) in an offline data store. Further, this data is a valuable source for determining the settings of thresholds to enable proactive management of customer experience.

It needs to implement powerful reporting tools, offering reports over selected time periods on historical data in predefined as well as customizable reports, including trend analysis, and which are readily available in the form of, for example, Web-based reports.

HP NGOSS solutions

Taking into account the concerns and challenges of CSPs we have touched upon earlier and the requirements of the solution resulting from the discussion above, HP is launching a set of innovative customer-centric assurance features and solutions across our portfolio:

- **HP Service Quality Management (SQM) Solution**

  The HP SQM solution helps CSPs optimize and consolidate their operations processes, shortening delays in service problem detection, identification, and resolution. It provides a comprehensive service quality management solution that allows service providers to easily manage services from end to end by monitoring associated service levels and quality of experience (QoE) in real time across the entire network and IT infrastructure.

  The HP SQM solution:

  – Provides early detection of customer experience problems through measuring, simulating user activity, and/or collecting usage data

  – Pro-actively identifies problems to initiate repair actions before they are noticed by end users through real-time monitoring and by correlating user experience metrics with network, service, and business-related data, while comparing against internal operations' service level objectives

  – Allows for prioritizing repair actions by comparing collected metrics against individual customer service level objectives and is easily configurable to notify relevant groups of people when target criteria are not met

  – Includes sophisticated capabilities that allow exploiting stored historical data for trend analysis to enable enhancements of customer perception over time and drive design of new service offerings

The HP SQM solution provides the foundation, necessary building blocks, and pre-designed models with smart KQIs to address the full spectrum of wire line and wire-less convergent service offerings. This includes network access technologies such as GSM, GPRS, 3G, cable, and xDSL. It enables all types of services like SMS, MMS, WAP, VoIP, IPTV, Internet, and e-mail, as well as complex value-added services (VAS) that are composed and delivered through a service delivery platform (SDP) in IP Multimedia Subsystem (IMS) architectures.

Figure 2. HP NGOSS—new features in customer assurance

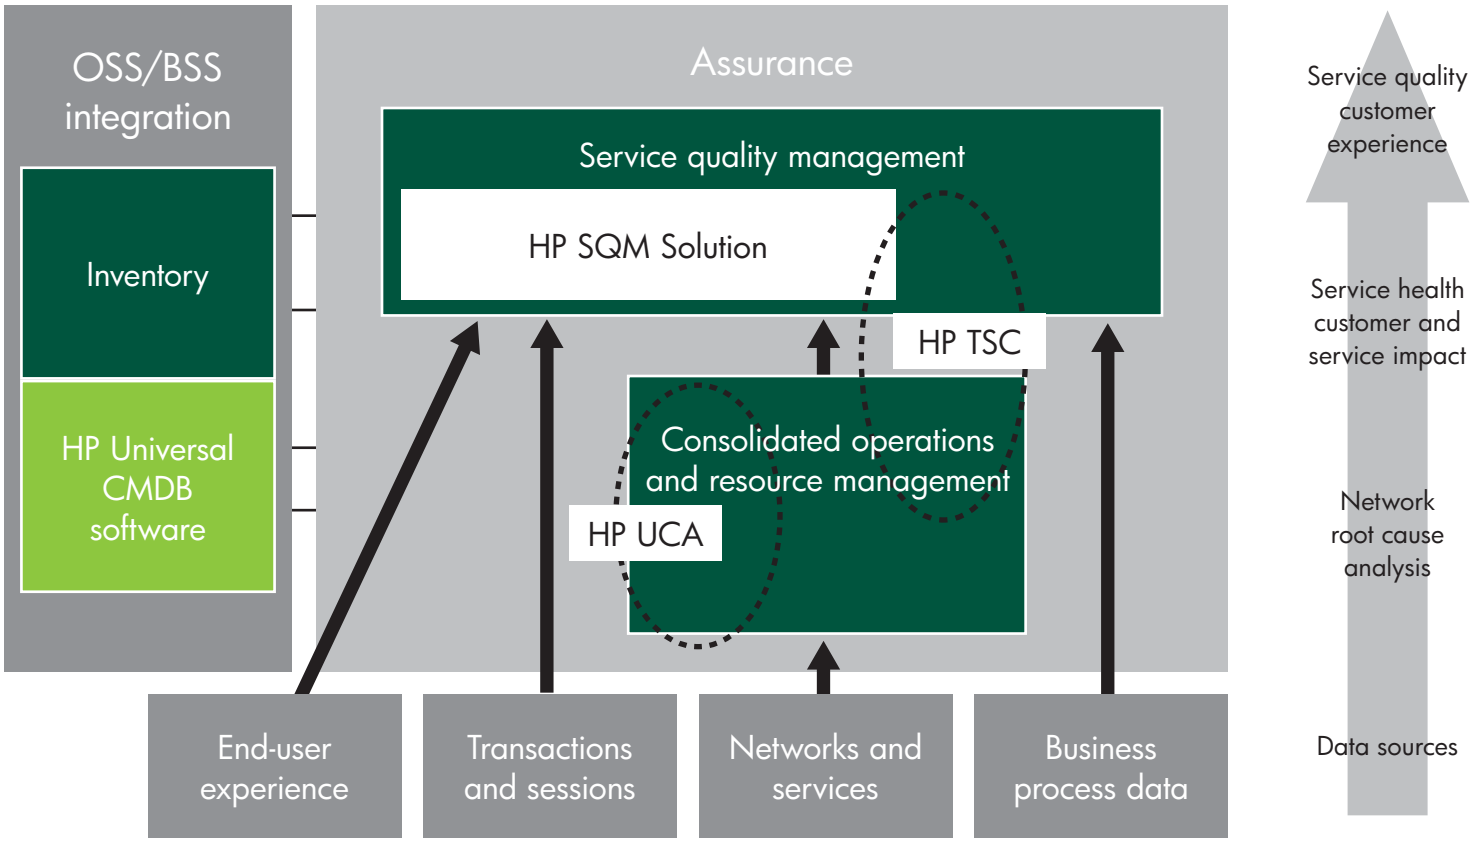

- **HP TeMIP Service Console (TSC)**

  HP TSC provides an end-to-end service view by monitoring the health of network resources and the services carried over the network. Built upon the HP TeMIP solution and the HP Service Management Foundation (see below), HP TSC uses the network resource data collected by HP TeMIP, continuously computes and propagates service component statuses in the service tree, and monitors compliance with service levels in real time.

  The HP TeMIP Service Console:

  – Includes state-of-the-art network problem detection and consolidation thanks to the HP carrier-proven manager of managers platform HP TeMIP

  – Leverages the innovative customer assurance capabilities of the HP SQM solution by implementing tight integration between the core SQM technology and the network resource management layer, resulting in resource-service-customer consolidation

  – Easily allows for extending its problem identification and resolution capabilities through off-the-shelf integration with HP UCA

- **HP Unified Correlation Analyzer (UCA)**

  HP Unified Correlation Analyzer (UCA) is a carrier-class problem determination, root-cause, and service-impact analysis product that has been designed to offer a radically different approach within large, complex, and fast-changing network environments. By utilizing a topology-based approach to correlation, HP UCA is able to accurately pinpoint the root cause and impact on services across multiple technology domains.

  The HP Unified Correlation Analyzer:

  – Quickly identifies problems through automated topology-based problem correlation and root-cause analysis across multiple network domains, and produces value-added service impact alarms to feed upper layers

  – Ideally complements manager of managers systems by automating major steps in the problem-resolution process thus reducing service outage time

  – Meets carrier-grade requirements through massive scalability and dynamic topology-model automatically kept in synchronization with network to avoid costly maintenance by operations staff

- **HP Service Management Foundation**

  The core HP SQM components included in both the HP SQM Solution and in the HP TeMIP Service Console are referred to as the HP Service Management Foundation.

  It is made up of a comprehensive library of pre-defined service components objects, service models, and KQIs based on the SID model. Once deployed, these models provide an extension to the UCMDB, referred to as the Telecom Universe.

  The Telecom Universe, together with the Service Designer tool, allows for quickly developing new service models in a standardized way, while providing great flexibility in the way pre-defined KQIs may be attached to service objectives with associated thresholds. The knowledge about the services, service components, and their interactions is captured in UML class diagrams, which are stored in the HP Universal Configuration Management Database (UCMDB) repository.

  The HP Service Management Foundation is comprised of:

  – Service modeling tools (Service Designer, KPI/KQI Modeler)

  – Telco Universe for the HP UCMDB (Universal Configuration Management Database)

  – HP BSM (Business Service Management) standard modules: dashboard, reporting service discovery, and synchronization, UCMDB

Conclusion

The new customer-centric assurance capabilities in HP NGOSS Solutions have been built from the considerations discussed in this white paper, while focusing on key areas for OSS improvement, and relating these to the needs encountered during the process used to address a typical customer experience problem.

Our hope is the new HP NGOSS solution capabilities will help CSPs move to an operations model where the customer perception drives their business strategy—and thus be able to transform their customers' experience.

Appendix

List of acronyms

BSS	Business support system
CEM	Customer experience management
CI	Configuration item
CMDB	Configuration management database
COTS	Commercial off-the-shelf
CSP	Communications service provider
DTV	Digital television
eTOM	Enhanced telecom operations map
GUI	Graphical user interface
ITIL	Information Technology Infrastructure Library
KPI	Key performance indicator
KQI	Key quality indicator
MDS	Mobile data services
NGOSS	New Generation Operations Systems and Software (TMF) Next-generation operational support systems (HP, common usage)
OSS	Operations support system
QoE	Quality of experience
RCA	Root-cause analysis
SID	Shared information/data model
SDP	Service delivery platform
SLA	Service level agreement
SLO	Service level objective
TCO	Total cost of ownership
TMF	TeleManagement Forum
UML	Unified modeling language
VPN	Virtual private network

References

GB921 TeleManagement Forum, Enhanced Telecom Operations Map (eTOM), The Business Process Framework, Release 7.1, January 2007.

GB922 TeleManagement Forum, Shared Information/Data (SID) Model, GB922, Release 6.0, November 2005.

Technology for better business outcomes

To learn more, visit www.hp.com/go/ngoss

© Copyright 2009 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

4AA2-5137ENW, March 2009