Measuring & Managing End-to-End Quality of Service (QoS) Provided by Linked Chains of Application and Communications Services

Dr. Arthur S. Robinson (art.robinson@stdc.com) and David Lounsbury (d.lounsbury@opengroup.org)

Abstract

This paper highlights performance and dependability problems facing organizations that are becoming increasingly dependent on the use of linked internal and external application and communications services to support their mission critical system requirements. It also summarizes a series of recent research advances that have the potential of providing solutions to this problem space by enabling coherent measurement and management of each link in these end to end chains, including the dynamic management of competing application and communications service requests. Many of these advances are the result of DARPA's ongoing Quorum QoS middleware research program, encompassing technologies that will support the specification, measurement and control of each link in a given critical service chain. When converted into industrial products and services, these technologies will enable the creation of metrics based specifications for each link in a given service chain, augmented by QoS Metrics Services (QMS) that will provide dynamic measurements of the actual services being provided. By comparing actual to specified services, link Resource Managers will then be able to dynamically allocate application and communications resources, as required to assure compliance with link specifications. Resource Managers, in turn, will be driven by dynamically assigned resource management policies. These policies will provide applicable current rules for allocating resources when requested services exceed available resource capabilities.

By the end of 2001 a substantial number of multi year, QoS focused, research results from projects funded under DARPA's Quorum Program will reach levels of maturity suitable for technology transfer into system services and products. The Quorum Program, conceived and directed by Dr. Gary Koob, correctly anticipated the growing need for assured end-to-end quality of service in computing and communications systems operating with shared pools of system resources. Positive results from Quorum research tasks are poised to make major contributions to organizations that are becoming increasingly dependent on the use of linked internal and external application and communications services to support their mission critical system requirements. Increasingly, these dependencies are evolving to include the majority of mission critical e-commerce and military computing and communications systems.

The concepts presented in this paper are the result of both participation in the Quorum program and related ongoing collaborative tasks being pursued by System/ Technology Development Corporation (http://www.stdc.com) and The Open Group (http://www.opengroup.org/RI/).

To place the overall end-to end QoS problem into perspective, it is clear that the emergence and rapid acceptance of Internet and Intranet technologies is providing commercial and military systems with the opportunity to conduct business at reduced costs and greatly increased scales by capitalizing on the use of shared-cost communications and computing infrastructures. However, to take advantage of this opportunity, organizations are becoming, "increasingly dependent on large-scale distributed systems that operate in unbounded network environments (IEEE *Internet Computing* 11/99)".

As the value of these transactions grow, companies are beginning to seek guarantees of dependability, performance, and efficiency from their distributed application and network service providers. To provide adequate levels of service to customers, companies eventually are going to need levels of assured operations similar to those received in the past from mainframe "Glass Houses". These capabilities include policy-based prioritization of applications and users competing for system resources; guarantees of levels of provided performance, security, availability, data integrity, and disaster recovery, and adaptivity to changing load and network conditions.

Typical Multiple Zone Computing and Wide Area Network (WAN) Service Links Required by Critical End To End Commercial and Military Applications

In a typical end to end application, a user request to zone A may trigger data feedback from multiple participating zones - say A and B. User selections from A and B results then could, typically, trigger data from additional zones - say C and D. User commands, based on data from sites A, B, C and D may then trigger:

- authentication request to zone E
- authorization request to zone F
- action request to zone G, with zone G then tracking and reporting on action progress

The net effect of these sequences of interactions is that end-to-end QoS results will depend on services received from each link in the chain. When services are provided on a "best efforts" basis, this will often result in intermittent critical end-to-end applications failures.

Figure 1, "Top Level Component Map for Federated E2E Solutions" was created by one of TOG's Open Systems QoS Task Force teams. It suggests the recursive nature of

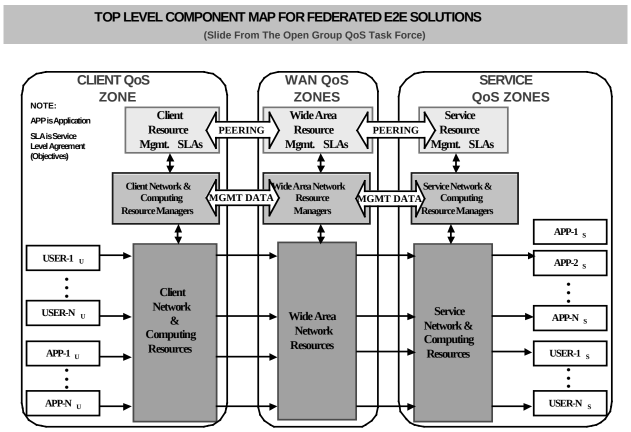

these computing and communications interactions. Typically, computing and localized communications interactions within a given zone will use WAN resources to enable multiple communications interactions with a series of remote service zones. As the figure suggests, both within zones, and over the WAN, mechanisms will be needed to manage competing resource requirements of multiple applications sharing these computing and communications infrastructures

Figure 2, "Example – Multiple Applications" suggests how multiple applications will be competing with one another for available computing and communications resources.

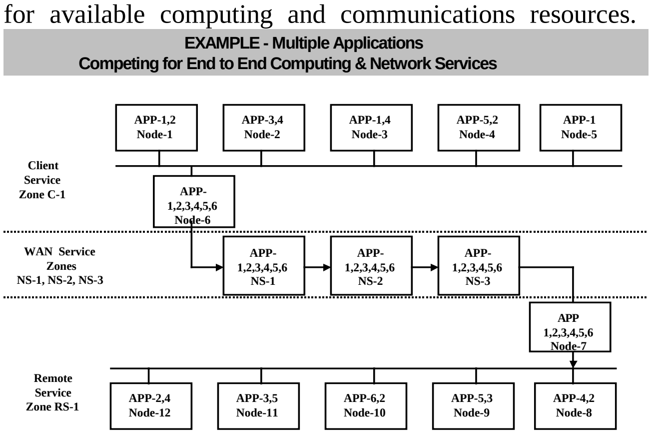

Clearly, to assure end-to-end QoS for any given application, means will have to be introduced for enabling policy based allocation of resources between these competing applications. By controlling relative allocations of computing and communications resources, policies (that can themselves be dynamically changed) will control the end-to-end QoS experienced by each of these applications. We will explore three dimensions of this end-to-end QoS assurance problem:

- Approaches to specifying, measuring and controlling each link in the chain, with the aim of allocating shared resources in accordance with dynamically specified resource allocation policies.
- Approaches to managing resources at multiple, cooperating, system resource levels, such as node level, LAN level and WAN level, and
- Approaches to augmenting each link in the chain with alternative resources, to be used when primary resources are not available, and/or to use primary and secondary resources to share system workload.

If we consider, first, a single link in the chain. To make it possible for that link to dynamically adapt to changing system conditions, we will clearly need the major functions suggested in Figure 3, "Specifying, Measuring and Controlling Each Link". These include means for

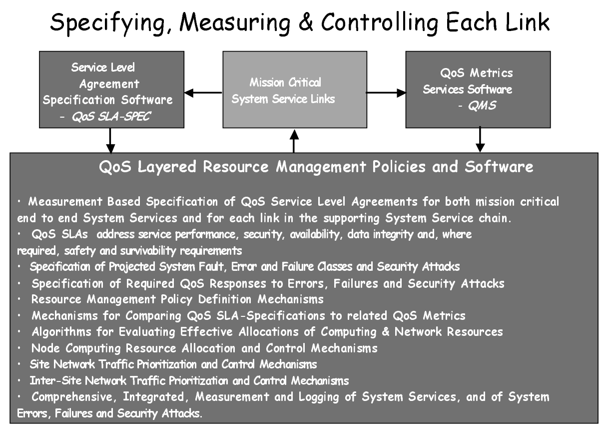

specifying, in terms that can be confirmed through measurements, the services that each link must provide; means for measuring the actual levels of service that are being delivered; means for comparing actual to specified levels of service; and means for determining and then controlling actions that drive actual delivered services to meet or exceed specified levels of service.

When reviewing each of the QoS Layered Resource Management Policies and Software outlined in Figure 3, it is important to note that an important part of the link specification will include clear specifications of the error and failure classes and the security attacks that the link has to withstand, and the acceptable levels of link responses to those stresses. Additional functions that QoS Layered Resource Management Policies and Software will have to provide will include:

- Resource Management Policy Definition Mechanisms

- Mechanisms for Comparing QoS SLA-Specifications to related QoS Metrics
- Algorithms for Evaluating Effective Allocations of Computing & Network Resources
- Node Computing Resource Allocation and Control Mechanisms
- Site Network Traffic Prioritization and Control Mechanisms
- Inter-Site Network Traffic Prioritization and Control Mechanisms
- Comprehensive, Integrated, Measurement and Logging of System Services, and of System Errors, Failures and Security Attacks.

Figure 4 presents a simplified overview of Basic Resource Management Specify, Measurement, Decide and Control functions. While they will, increasingly, be automated, they should be recognized as normal functional practice when humans manage system resources.

SLA Driven Resource Management Functions

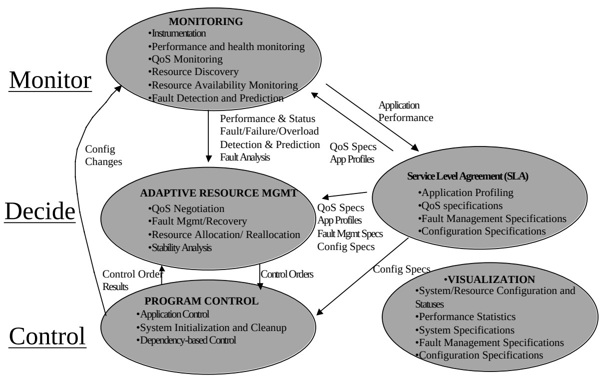

Figure 5, "Resource Management Levels", highlights the need for managing resources at multiple, cooperating, system resource levels. In this example, cooperative, multi level resource management functionality operates at the Node, LAN and WAN levels.

Resource Management Levels

- Accepts directives from higher levels
- Provides status to higher levels
- Manages lower levels
- Receives information about performance of lower levels

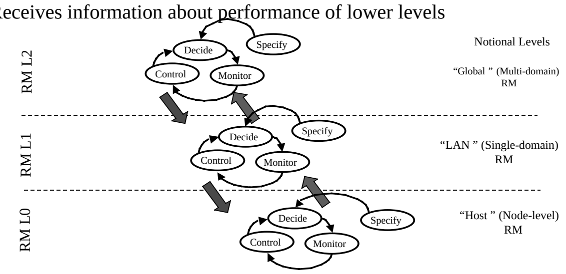

The third dimension of assuring end-to-end system QoS functionality relates to approaches that augment each link in the chain with alternative resources, to be used when primary resources are not available, and/or to use both primary and secondary resources to dynamically share current system workload. This approach is suggested by figure 6, below:

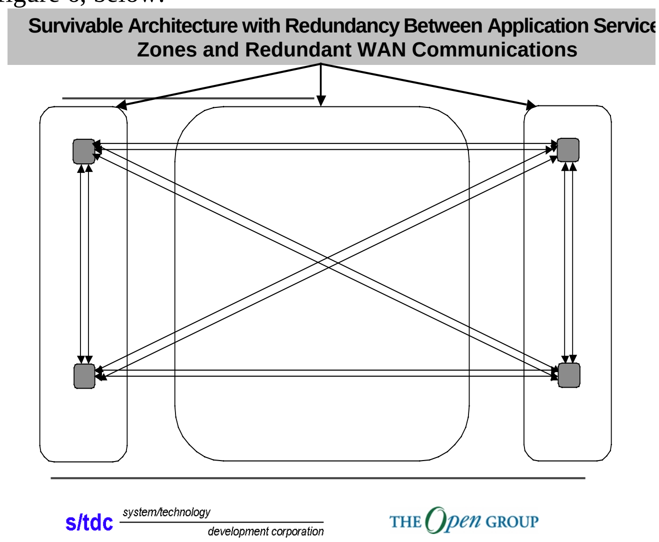

While appropriate levels of redundancy can significantly improve end-to-end application dependability, they usually will also involve significant increases in both cost and design complexity. For a given system application, cost trade offs will have to consider both these factors and the cost impacts of the lower levels of dependability that would be provided by designs that do not provide fault tolerance capabilities in their computing and communications links.

Related Quorum Research Results.

Page constraints limit the level of detail that can be included in this paper concerning the many evolving Quorum research results that have the potential for use in the resolution of each of the end-to-end QoS measurement and management areas that have been described.

When the Quorum Program ends later this year, S/TDC and The Open Group are planning to create detailed on line resources with descriptions and pointers to the full range of related research documentation. Interested parties will be able to access its contents at **http://www.stdc/QoSResearch.** Further, The Open Group has started an end-to-end QoS initiative that is addressing both impacts of Quorum Program research advances and related ongoing QoS related product and service developments. Interested parties can join this activity by contacting **http://www.opengroup.org/qos/**

Quorum Research References

Globus
http://www.globus.org
The Globus project is developing fundamental technologies needed to build computational grids. Grids are persistent environments that enable software applications to integrate instruments, displays, computational and information resources that are managed by diverse organizations in widespread locations.

AquA - Adaptive Quality of Service for Availability
http://www.dist-systems.bbn.com/projects/AQuA
The objective of the AquaA project is to make orders of magnitude improvements in the specification of availability requirements, in the prediction of availability under varying conditions, and in the adaptability of the applications and resource management strategies by providing infrastructure mechanisms to support dynamic behavior. The goal is to extend the CORBA based Application Program Interface (API) for specifying an application's availability requirements through defining a set of acceptable operating regions, and adapting when various thresholds of predicted availability will not be met.

Distribute Object Integration for the Quorum Program
http://www.dist-systems.bbn.com/projects/QuOIN
The objectives of the QuOIN project are to develop an adaptive middleware layer that is applicable to and integrates with a number of Quality of Service (QoS) oriented attributes, and to support transition of that environment to DoD and commercial interests. Quorum is currently supporting a number of technology initiatives, each addressing a single aspect of QoS (e.g. dependability, real time performance, security, as well as resource management mechanisms). This project is intended to unify a number of these Quorum technologies under a common distributed object based structure as an integrated concept, and to demonstrate, test, evaluate and deliver the composite result, as well as subset results, for technology evaluation, transition and adoption.

WSOA: Quorum Technologies Extensions Program
http://www.darpa.mil/ito/psum1999/H717
The WSOA program is focused on developing open systems architecture (OSA) and Quorum technologies that can be applied to directly address limitations of deployed embedded mission systems and communications technology including:

1. static resource management;
2. finite compute resources of airborne elements balanced against hard real-time constraints associated with normal mission activities; and
3. bandwidth of current tactical data links which constrain the operational utility of systems for collaborative planning and attack of mobile targets.

We are developing pilot vehicle interfaces for incorporating collaborative planning and browser capability into strike platform and C2 nodes incorporating Quorum software into F-15 and airborne C2 nodes. As a final examination, we will provide a live flight demonstration that will illustrate the significant gains in war-fighting capability provided by enabling technologies of collaborative planning, information mining, and embedded resource management.

Darwin: Resource Management in Application-Aware Networks
http://www.cs.cmu.edu/~darwin
Darwin research supports resource allocation along three dimensions: resource allocation in the "space" consisting of the physical network infrastructure and attached processing and storage resources; hierarchical resource management by different organizational entities sharing the infrastructure; and application-specific adaptation on multiple time scales, ranging from application startup to packet and cell scheduling. In all three dimensions, the mechanisms provide for extensive tailoring to application requirements.

Amaranth: Adaptive Multidimensional QoS
http://www.ices.cmu.edu/
The Amaranth project provides multi-dimensional, adaptive, assured Quality of Service (QoS) for heterogeneous distributed computing systems. In particular, Amaranth provides probabilistic guarantees of service encompassing real-time deadlines, dependability, cryptographic security, and application-specific performance. Amaranth incorporates adaptive, distributed optimization techniques to maximize system performance even in degraded operating modes while maximizing QoS stability. User-specific tradeoffs of service commitment duration, commitment firmness, and service dimensions provide appropriate service to all users even in a resource-limited system. A specific goal is to provide an order of magnitude improvement in system efficiency over static resource allocation, while assuring a comparable QoS.

Resource-Centric Microkernel Communication Services
http://www.cs.cmu.edu/~rajkumar/resource-kernel.html
A resource kernel is defined as mechanisms that provide timely, guaranteed and protected access to system resources. The resource kernel allows applications to specify only their resource demands, leaving the kernel to satisfy those demands using hidden resource management schemes. This separation of resource specification from resource management allows OS-subsystem-specific customization by extending, optimizing or even replacing resource management schemes. As a result, this resource-centric approach can be implemented with any of several different resource management schemes. The resource

kernel gets its name from its resource-centricity and its ability to apply a uniform resource model for dynamic sharing of different resource types, take resource usage specifications from applications, guarantee resource allocations at admission time, schedule contending activities on a resource based on a well-defined scheme, and ensure timeliness by dynamically monitoring and enforcing actual resource usage. Generally, a QoS manager sitting on top of a resource kernel can make adaptive adjustments to resources allocated to applications.

Remulac: Resource management under language and application control

http://www.cs.cmu.edu/~cmcl/

Remulac's research goal is to simplify the design and implementation of applications that are network-aware (i.e., adjust their resource demands to resource availability or reserve the required resources). Network-aware applications enable network-wide end-to-end resource management. The Remulac (REsource Management Under Language and Application Control) project develops a uniform method to provide resource management for applications so that they can benefit from the availability, replication, and cost-effectiveness of networked environments.

RT-ARM Real-Time Adaptation

http://www.htc.honeywell.com/projects/arm

RT-ARM's goal is to develop techniques for adaptive resource management in dynamic environment that provide predictable real-time performance needed to meet mission-level objectives. RT-ARM focuses on providing dynamic adaptations as a key ability through which systems can continue to meet functional and performance requirements in the face of changes**.**

MSHN: Management System for Heterogeneous Networks

http://cisr.nps.navy.mil/Research/mshn/index.html

The objective of the MSHN project is to explore the application of adaptive and heuristic matching and scheduling techniques and modern distributed security methods to a distributed RMS which allows heterogeneous resources to be accessed by both MSHN-controlled and external applications.

Quorum Technology for Shipboard Computing

http://www.nswc.navy.mil

This project focuses on the evaluation and validation of advanced computing architecture concepts and advanced computing technology components. The primary unifying theme of both architecture concepts and technology components is providing the quality-of-service (QoS) needed to carry out Navy missions in the 21st century. These architecture concepts and technology components are evaluated for suitability and transition for Navy real-time, mission critical advanced computing applications.

Computing Communities: Information Survivability via Adaptable Virtualization

http://www.cs.nyu.edu/cc

The goal of the Computing Communities project is to design, develop, and implement a software system architecture and supporting mechanisms for integrating autonomous computing and information resources into Computing Communities (CC for short). A given CC will provide applications with a unified execution platform, which is dynamically changeable and automatically configurable and supports features such as distribution, migratability, global scheduling, and adaptation to changes in resource availability.

Microfeedback for Adaptive Resource Management

http://www.cse.ogi.edu/DISC/projects/quasar

The QUASAR project is developing the key concepts and abstractions for systemic quality of service adaptation, and prototyping them in real system components. Key concepts and abstractions include a general model for QoS adaptation, microlanguages for declarative specification of QoS requirements, meta-interfaces for communication of QoS-related events, and the microfeedback abstraction for building adaptive components with well-defined QoS behavior.

DeSiDeRaTa: QoS Management Tools

http://desidrta.uta.edu/

Mission-critical real-time software systems frequently operate in dynamic environments, thereby precluding accurate characterization of the applications' workloads by static models. Thus, guarantees of real-time QoS based on a priori characterizations are not possible. A posteriori approaches offer significant benefits in such contexts, including the ability to function correctly in dynamic environments (through adaptability to unforeseen conditions) and higher actual utilization of computing resources. Thus, the DeSiDeRaTa project is developing adaptive resource management middleware algorithms for distributed real-time systems that must operate in dynamic environments (contexts such as war-fighting and the internet).

Cactus: An Integrated Framework for Dynamic Fine-Grain QoS

http://www.cs.arizona.edu/cactus/

The objective of the Cactus project is to develop a design and implementation framework for distributed service supporting customizable dynamically adaptive fine-grain Quality of Service (QoS) attributes related to dependability, real time, security, performance, consistency, and other quality attributes in distributed systems.

WAFT

http://www.cs.ucsd.edu/users/marzullo/WAFT

The objective of the WAFT project is to make fault-tolerance a practical feature that can be added to the next generation of distributed applications. We have identified three objectives that are necessary to realize our vision:

1.Fault-tolerance techniques must make as light demands on the applications as possible. Current successful fault-tolerance techniques, such as the state machine approach (which has been embodied in toolkits like ISIS), make heavy demands and have only been worth the trouble for mission critical applications like HiPer-D and the Boeing 777.

2.Fault-tolerance techniques must be scalable and usable in a wide-area environment. Approaches, like ISIS, that have a well-developed methodology (based on the State Machine approach) do not scale well and have proven hard to apply to wide-area environments.

3.Fault-tolerance techniques must be applicable to the constraints of modern Internet usage. The lack of global trust and the possibility of cyber attack must be addressed.

End-to-End QoS and Resource Management
http://epiq.cs.uiuc.edu

The EPIQ project is developing a framework within which end-to-end quality-of-service (QoS) and resource management strategies can be customized and integrated to provide guaranteed services of negotiated quality to time-critical C3I applications. Specifically, the framework provides end-to-end QoS management to each flexible application system and enables it to adapt the quality of its services to dynamic changes in requirements, demands for resources and availability of resources.

Quality of Service Dynamic Validation
http://www.cs.uoregon.edu/research/qos/

The ASSERT project seeks to develop an integrated approach to both pre-deployment and post-deployment validation. To support scalability, transferability to DoD systems, and interoperability with other Quorum technology, technical objectives for the ASSERT approach include: a common, industrial-strength method for modeling both functional and QoS requirements, end-to-end automated support specifying and checking requirements, capabilities for generating QoS contract middleware from specifications, and real-time QoS monitoring and trending. ASSERT's dynamic monitoring and trending capabilities are extended to provide command-level C3 decision support for new or evolving missions.

TAO (The ACE Orb)
http://www.cs.wustl.edu/~schmidt/TAO.html

This project is aimed at the design, prototype implementation, and demonstration of a high performance distributed object environment that supports end-to-end quality of service (QoS) guarantees. In particular we are developing a CORBA compliant, real-time object request broker (ORB) end system that provides deterministic and statistical QoS guarantees as well as the traditional best effort service. This real-time distributed object environment is essential for harvesting the potential of various technologies, such as object-oriented middleware, high-speed networks that can support QoS guarantees, and multimedia devices for supporting many mission critical applications.

QMS - QoS Metrics Services
http://www.stdc.com/QMS

QMS provides a message oriented transport mechanism and a set of services that facilitates communication between QoS data consumers and data producers. QMS contains a cross platform, robust, data agnostic publish/subscribe framework for this communication. A design goal of QMS citizens has been that they should themselves be well behaved, QoS aware, adaptive applications designed to fit well into a distributed QoS sensitive environment. Rapid failure detection and inbuilt security are other high priority design considerations. QMS uses XML for data representation, as well as for configuration and reports.