



*MEADEP White Paper*

---

# Availability Measurement, Modeling and Prediction for Mission Critical Systems

April, 2000

WP2000-1

SoHaR Incorporated  
8421 Wilshire Blvd., Suite 201  
Beverly Hills, CA 90211  
323-653-4717

[www.sohar.com/meadep](http://www.sohar.com/meadep)  
[info@sohar.com](mailto:info@sohar.com)

---

## Table of Contents

Introduction .....	1
Information System Availability Modeling and Prediction.....	2
Calculating Availability.....	3
Valuing Service Level Agreements .....	4
Value of a Restoration Time Service Guarantee.....	4
Valuing Additional Capacity.....	5
Conclusions .....	6
About SoHaR Incorporated.....	6

## List of Acronyms

CIO	Chief Information Officer
DBMS	Database management system
ISP	Internet Service Provider
MEADEP	Measurement-based Dependability software
MTBO	Mean time between outages
MTTF	Mean time to failure
MTTR	Mean time to repair/restore
RAID	Redundant Array of Inexpensive Disks
WAN	Wide Area Network

---

© Copyright 2000 by SoHaR Incorporated. All rights reserved

Windows NT is a Trademark of Microsoft Corporation  
Sprint is a Trademark of U.S. Sprint  
Uunet is a Trademark of MCI Worldcom  
eBay is a Trademark of ebay.com  
Cisco is a Trademark of Cisco Systems Inc.  
Compaq is a Trademark of Compaq Computer Corp.  
Solaris and Starfire are Trademarks of Sun Microsystems Inc.

## Summary

*A CIO or system administrator responsible for enterprise-level decisions on maintaining and increasing uptime needs modeling and prediction to make strategic decisions. Network monitoring services or software can provide current or past history on individual applications, servers, or networks. However, this information can not address higher level questions such as the combined availability of multiple systems performing a mission critical enterprise function (e.g., order entry, fulfillment, and CRM), the value of an availability guarantee in a service level agreement, or allocation of the budget to capacity vs. redundancy. The paper demonstrates how modeling and prediction using the MEADep tool developed by SoHaR can be used to make strategic decisions.*

---

## Introduction

In a recent article, a major Information Technology trade publication, put it bluntly: "It's now beyond discussion: Large e-business operations simply must be up all the time."<sup>1</sup> In the words of the Chief Information Officer (CIO) of one of the 10 largest traffic web sites, "Availability is as important as breathing in and out [is] to human beings."<sup>2</sup> Mission critical information system outages now regularly make the front pages of major general circulation newspapers<sup>3</sup>.

While the problem is well known, good solutions for analyzing and solving system level availability problems are not. Network performance monitoring software is often looked to for answers to availability issues. However, the primary function of such software is to identify problems requiring immediate attention. When used for availability questions, the sheer volumes of data gathered from performance monitors has stifled many an IT effort to analyze availability and multi-system questions. At the other extreme are Internet-based services that continuously assess the performance and availability of web sites. The information provided by these services is external and therefore not sufficiently detailed to provide insight into the problem areas of "back-end" systems. Neither network performance monitors nor external web-site monitoring can answer strategic questions such as

- My catalog sales databases are down too often. What architectural changes do I need to increase reliability?
- How much redundancy and what service provider guarantees do I need in order to support a given transaction level with an availability of greater than 99.99%?
- I may know which networks, backbones, or servers are failing most often, but which are the biggest contributors to the downtime of my mission critical information services?
- How can I measure the combined availability and performance benefit of an IT investment?
- How can my budget be spent most effectively to increase availability and performance?
- What are the highest impact items to negotiate in my service level agreement?

---

<sup>1</sup> Richard Winter, "The E-Scalability Challenge", *Intelligent Enterprise*, December 21, 1999, Volume 2 - Number 18

<sup>2</sup> Kal Raman, Sr. VP and CIO of drugstore.com as quoted by Susan E. Fisher, "E-business redefines infrastructure needs", *Infoworld*, January 7, 2000, available from [www.infoworld.com](http://www.infoworld.com)

<sup>3</sup> See, for example, Gary Strauss, "When Computers Fail", *USA Today*, page 1A, December 7, 1999

- Is the benefit of replacing a single server with a cluster and/or a RAID subsystem worth the acquisition and ongoing support cost?

This paper introduces a tool called MEADEP (Measurement-based Dependability) that will provide the answers through modeling, prediction, and measurement. MEADEP can be used during the architectural phase to assess the combined availability and performance of system alternatives. During operations, it can be used to assess system and enterprise functional-level availability and to identify the best opportunities for improvement. Through its graphical user interface, the tool can be easily learned and immediately applied by busy professionals. Its hierarchical modeling approach enables large information systems to be represented accurately and completely. Its sophisticated analytical methods provide consistent and repeatable results for high availability systems. MEADEP's data analysis and trending functions enable users to easily measure the availability of individual servers and workstations to assess the overall network or system availability.

## Information System Availability Modeling and Prediction

In order to show the benefits of MEADEP, we present case studies on analyzing contributors to downtime, analyzing a service level agreement, and assessing availability from the perspective of uptime and performance, and measurement of results. Our examples will be based on a simplified description of the eBay system as reported in the Los Angeles Times<sup>4,5</sup>. Figure 1 shows the system architecture, consisting of dual redundant backbone networks, routers, front-end web servers, and third tier Database Management System (DBMS) servers.

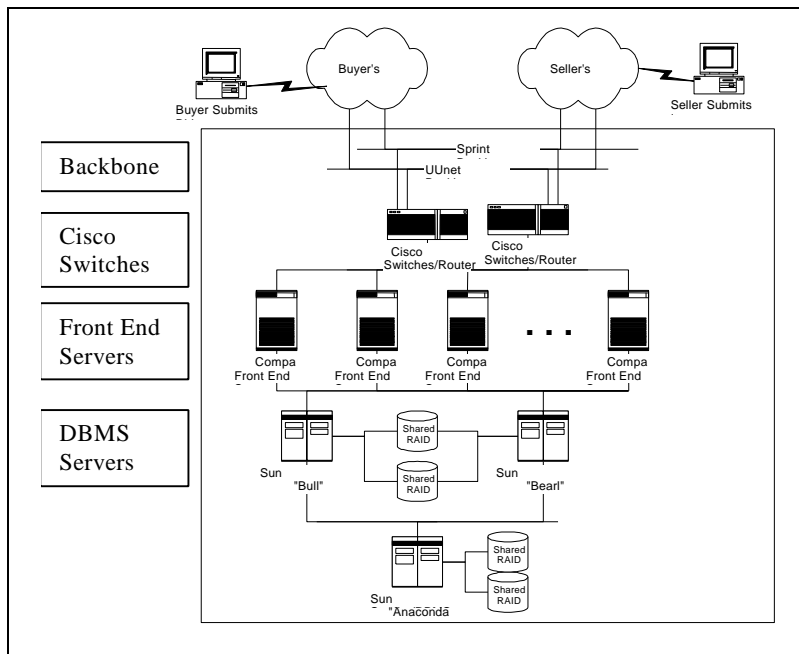



Figure 1. Simplified Diagram of Large e-Commerce Site Based on e-Bay

<sup>4</sup> Joseph Menn, "Prevention of Online Crashes in No Easy Fix", *Los Angeles Times*, October 16, 1999, Section C, Page 1

<sup>5</sup> A separate application note describing the details of the model is available from SoHaR either through the web site ([www.sohar.com/meadep](http://www.sohar.com/meadep)) or by contacting us directly

## Calculating Availability

The eBay system described above was modeled in MEADEP using a total of 9 hierarchical models as shown in Figure 2. The Figure also shows examples of the three lower level models. An experienced MEADEP user could create this or a similar model in less than 3 hours using the MEADEP graphical user interface. We used data from existing systems for failure rates and recovery times of all of the components. MEADEP allows the utilization of data from system event log files as well.



### Examples of Lower Level Models (2 of 9)

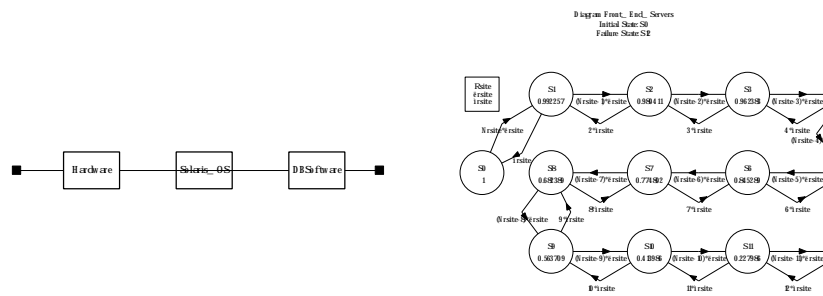


Figure 2. Portions of the System Reliability Model

The predicted availability is 99.6%, corresponding to about 35 hours per year of downtime. However, the real value of the model is what it can show about the underlying system. Figure 2 demonstrates one of the insights we can gain by showing which of the major network subsystems are failing the most: in this case, the front-end server subsystem. However, as lower level models would show, it is actually quite expensive to decrease the failure rate of the server subsystem. A more cost-effective measure would be to improve the reliability of the backbones through enhanced service level agreements. The following section provides an example of how MEADEP can be used to analyze such a service level agreement.




Figure 3. Failure Rates of Major Subsystems

## Valuing Service Level Agreements

Uptime translates directly into transactions and then into revenue for the enterprise. In this section, we'll show how the MEADep can be used to value restoration time and reliability service guarantees, two of the key elements of a service level agreement (SLA)<sup>6</sup>. According to a recent survey<sup>7</sup>, the average WAN downtime was 67 minutes., and that outages occurred 23 times per year, or approximately once every 380 hours. We'll assume that both of the service providers in our example have offered somewhat better service guarantees: a 1 hour restoration time and a 500 hour MTBO (corresponding to an availability of 99.8%). For the purposes of valuing uptime, we'll assume the reported 650 transactions per minute at the eBay site and assume a net revenue of \$0.10 per transaction.

### Value of a Restoration Time Service Guarantee

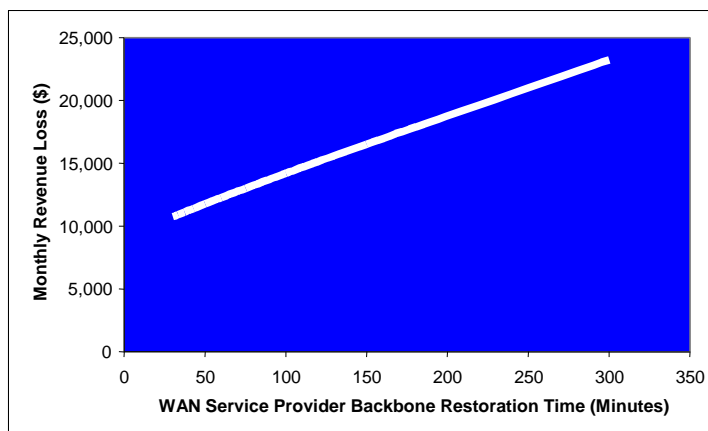


Figure 4. Impact of Service Provider Restoration Time on Annual Lost Transactions

We'll first discuss the value of a lower service restoration time (sometimes called a downtime) guarantee. Figure 4 shows the impact of service provider restoration time on revenue using MEADep's parametric analysis function<sup>8</sup>. The results show a monthly loss of about \$3,000 for each additional hour that it takes the service provider to restore the network after an outage. For a

restoration time of 1 hour, there will be about 115,000 lost transactions per month, or \$11,500 per month in lost revenue. If the downtime is decreased by 30 minutes, the lost revenue decreases to about \$10,000 per month. Thus, the value of this higher service level guarantee is about \$18,000 per year.

### Value of an MTBO Service Guarantee

Figure 5 shows the impact of the MTBO of an Internet backbone on the number of lost transactions (the MTBO of the second Internet backbone is held constant). As the MTBO is increased from 100 to 1000 hours, the revenue increased by approximately \$25,000 monthly. However, Unlike the restoration time shown in Figure 4, there are decreasing returns for increased MTBOs, and we may be able to take advantage of this to reduce costs or reallocate our budget. For example, if we are willing to accept a lower service level guarantee of 380 hours MTBO (the average achieved in the Infonetics report), a 120 hour MTBO decrease relative to our baseline 500 hour service guarantee, the average monthly revenue loss will be about \$900. On the other hand, if we wanted to increase our MTBO by the same 120 hours, the monthly revenue increase would only be about \$550.

<sup>6</sup> Although availability is the focus of this example, it is not necessarily always the most important attribute of uptime. For example, if a system is taken down in periodically in accordance with a scheduled maintenance plan, then the reliability, i.e., the probability that the system will be up during the maintenance interval is the primary figure of merit.

<sup>7</sup> "WAN Downtime and SLAs", Infonetics Research, San Jose, CA, December, 1998

<sup>8</sup> We have assumed a transaction rate of 650 per minute, Mean Time Between Outages of 500 hours for both backbone networks and a restoration time of 1 hour for the service provider of the second network, in this case, assumed to be UUnet.

It must be emphasized that these results are with the presence of *two independent and parallel* service providers. If a single backbone is used, the relative importance of MTBO and outage restoration time will differ.

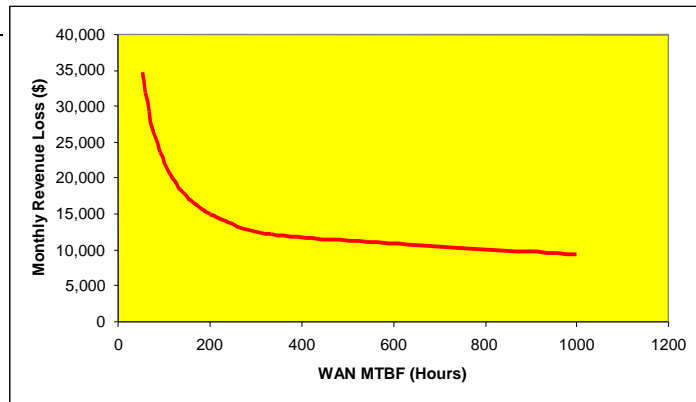


Figure 5. Impact of Backbone Outage Rate on Web Site Downtime

### Valuing Additional Capacity

At first glance, the value of additional capacity might look easy to determine: the greater the capacity, the greater the number of transactions that can be processed. However, when we consider reliability and availability, the considerations can become significantly more complex. For example, our eBay site model has two independent backbones, but neither has the capacity to carry the full traffic load (the model assumes that each has half the capacity). Given the MTBOs and restoration times of the networks, MEADEP can be used to evaluate the value the additional capacity through its *reward* function. Figure 6 compares the monthly revenue loss for two alternatives:

- *Half capacity backbones:* 2 WAN backbones each with half the capacity required for the average load, and
- *Full capacity backbones:* 2 WAN backbones each with the capacity for the entire load

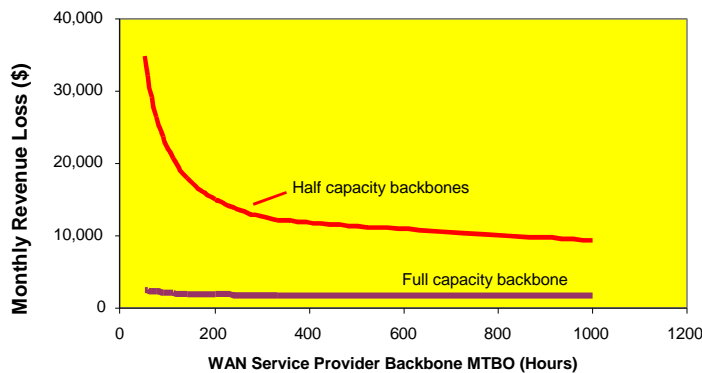


Figure 6. Revenue losses due to outages of WAN backbones

The results show that increasing the capacity of both backbone WANs so that either could handle the full load could decrease the monthly revenue loss by as much as \$30,000 or as little as \$9,000 depending on the reliability (as measured by mean time between outages) of the service providers<sup>9</sup>. The lower the reliability (MTBO), the greater the benefit of increasing the capacity.

With these results, it is also possible to assess the value of a service level agreement with the additional dimension of capacity. For example, it may be that provisioning each of the backbones with double the capacity is a lower cost option than securing service level agreements with both WAN backbone service providers. With the MEADEP reward function, it is also possible to consider other options such as provisioning the first backbone with full capacity and the second with half capacity, or provisioning each with 75% of the capacity. The optimum configuration depends on a number of factors including the traffic profile and the MTBO and restoration times of the WAN backbones, the load profile of the server, the value of the traffic, and other system-specific factors.

<sup>9</sup> Under the assumption of a uniform 650 transactions per minute and a value of \$0.10 per transaction outline, the 1 hour response time, and the 500 hour MTBO on the second backbone WAN discussed above.

## Conclusions

Modeling is an effective tool to analyzing on-line systems and developing the answers to strategic questions to achieve higher availability and performance. The methodology described here can be applied not to other computer architectures and configurations. For example, published MEADEP-based analyses have included air traffic control systems and nuclear reactors. These studies are available from the SoHaR website, at [www.sohar.com](http://www.sohar.com). Other applications of MEADEP have been in telecommunications, process control, and HVAC.

This paper has addressed the modeling capability of the MEADEP tool but not the data analysis (such as shown in Figure 7 at right). A separate paper will discuss the Data Collection System product which collects data from Windows and UNIX event logs and incorporates them into reliability models. More information on the data evaluation and analysis capability of MEADEP is available on the SoHaR web site.

## About SoHaR Incorporated

SoHaR is a consulting, research, and development organization which specializes in computing for critical applications (hence the name: Software and Hardware Reliability). SoHaR has participated in the development of satellite, missile, and aircraft control systems; ground-based transportation control; and nuclear reactor safety and control. Our areas of research include distributed systems fault tolerance; Internet-based logistics and maintainability applications; safety software verification and validation; and advanced software analysis tools for reliability analysis, sneak circuit analysis, and critical software testing. Since its founding in 1978, employees of the firm have authored more than 200 articles in refereed publications and conferences. In 1990, we received the U.S. Small Business Administration Prime Contractor of the Year award in Region IX (Southwest Pacific and Hawaii).

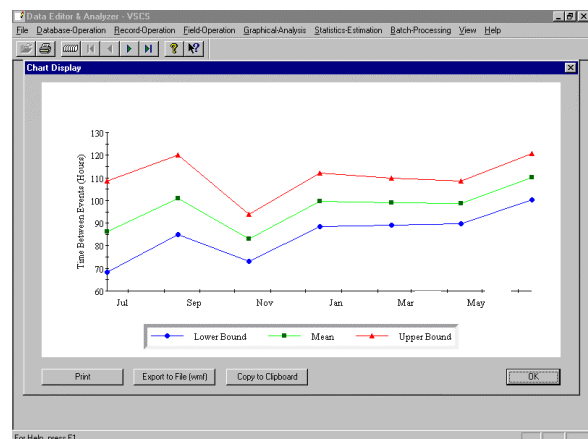


Figure 7. Time Trend for Mean Time Between Events