

Ensuring Successful Application Delivery

Jim Metzler
jim@ashtonmetzler.com

Goals and Non-Goals

- Goals

- To review the application delivery track
- To describe the current approach to application delivery
- To discuss the factors that are causing application delivery to become more difficult
- To describe a framework for successful application delivery

- Non-Goal

- To be definitive – organizations are unique and this is a broad, complex topic!

Agenda

- The Application Delivery Track
- The Current Approach to Application Delivery
- Factors that Make Application Delivery Difficult
- A Framework for Application Delivery
- Suggestions and Summary



The Application Delivery Track

- Ensuring Successful Application Delivery
 - Tuesday, May 22, 2007, 10:15 to 11:15
- Planning for Application Deployment
 - Tuesday, May 22, 2007, 11:30 to 12:30
- Services Oriented Architecture (SOA) and Web Services: A Reality Check
 - Tuesday, May 22, 2007, 2:15 to 3:15
- XML Gateways: Addressing the Security and Performance Impact of XML on the Network
 - Tuesday, May 22, 2007, 3:30 to 4:30

The Application Delivery Track

- Managing and Controlling Application Performance
 - Wednesday, May 23, 2007, 10:15 to 11:15
- Adding More Intelligence to the Network
 - Wednesday, May 23, 2007, 11:30 to 12:30
- Securing and Managing a Web services Environment
 - Wednesday, May 23, 2007, 2:15 to 3:15

Agenda

- The Application Delivery Track
- The Current Approach to Application Delivery
- Factors that Make Application Delivery Difficult
- A Framework for Application Delivery
- Suggestions and Summary



The Current Approach to Application Delivery

- In the vast majority of situations application degradation is noticed first by the end user.
- When the IT organization is made aware of application degradation, it often is not sure if indeed the application has experienced degradation.
- In most cases, it is assumed that the network is at fault and there is the temptation to *throw bandwidth* at the problem.
- Few IT organizations have a good handle on the flow of information inside of an application and hence can not evaluate alternative solutions.

Example

- An IT organization hosts an application in Atlanta, GA.
- Users in Hong Kong complain of poor performance.
- The IT organization confirms that indeed there is intolerable delay.
- The IT organization determines that if it redesigns the WAN it can make a modest improvement to the performance of the application.
- They have no idea if caching, compression, protocol acceleration will help.
- They are tempted to host a copy of the application in Hong Kong but have no idea about the traffic between the application server and the database server.

Common Application Delivery Myth

- Myth:

- If the network is running well (i.e., low delay, jitter, packet loss) then the application is running well.

- Reality:

- The network could be running at high utilization rates and the application is running well if the application is email or bulk file transfer.
- The network could be running at very low levels of utilization and have the application perform very badly if it is a chatty application running over a network with non-trivial delay.

Example

- Assume that it takes 200 application turns to complete a given transaction.
- On a LAN with 5 microseconds of delay, this takes 1 millisecond, which is not noticeable.
- On a WAN with 100 milliseconds of delay, this takes 20 seconds.
- If you throw bandwidth at this problem, it will not reduce the 20 second delay.
- If you throw bandwidth at this problem, it will increase cost.


Jim Metzler's Premises

- If you work in IT, you either develop applications or you deliver applications.
- Successful application delivery requires a top down approach that focuses on the application and the user's experience using that application.



What is Application Delivery?

In most of the articles in the trade journals, application delivery refers just to techniques to optimize the network and applications that transit the network.



Graphic source: http://www.pixelmonger.com/art_flavors.html

Application Delivery Framework

- There are four components of Application Delivery. They are:
 - Planning
 - Network and Application Optimization
 - Management
 - Control

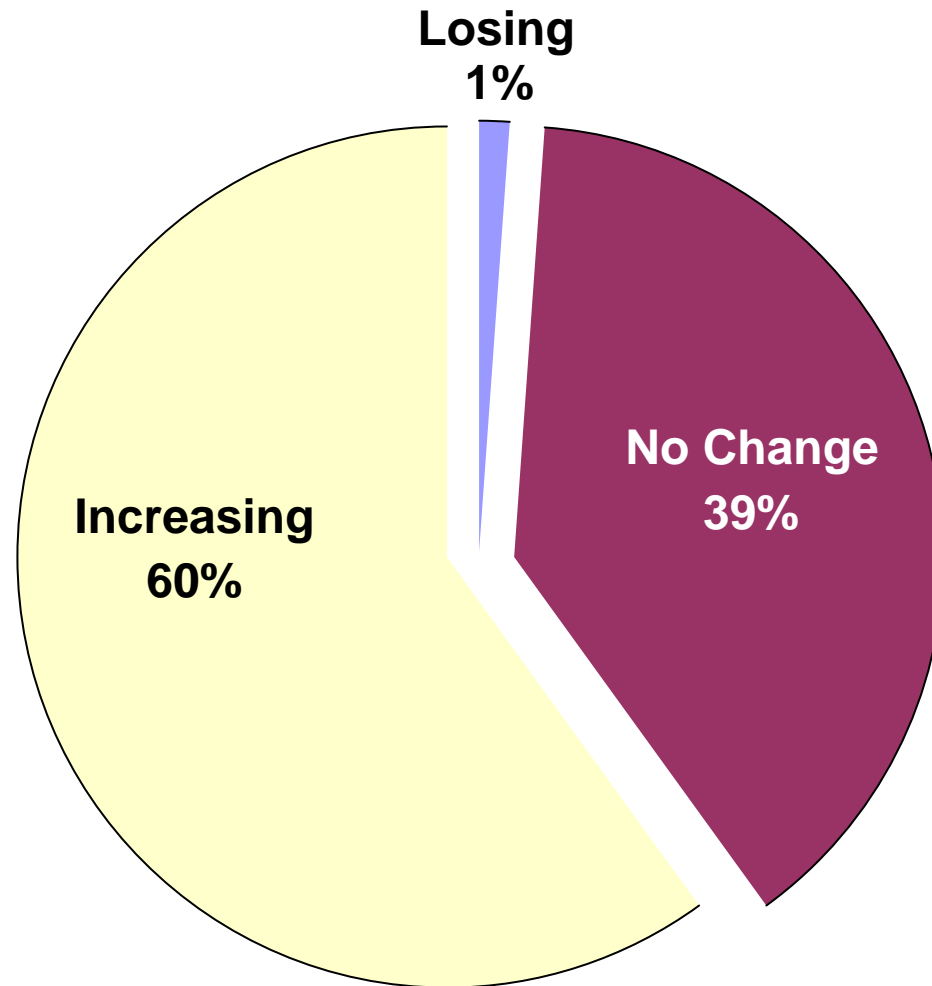


Agenda

- The Application Delivery Track
- The Current Approach to Application Delivery
- Factors that Make Application Delivery Difficult
- A Framework for Application Delivery
- Suggestions and Summary



The Change in the Importance of Managing Application Performance



The Factors That Make This Important and Difficult

- Repeat: In the vast majority of times when an application degrades, it is noticed first by the end user.
- Business unit managers do not care significantly about the IT infrastructure.
- Business critical business processes increasingly depend on IT. Example: Wal-Mart and SCM.
- Applications do not typically degrade on a daily basis, but when certain ones do it can have a significant impact. Example: High Tech and quarter close.

The Factors That Make This Important and Difficult

- During application development, few application development groups place much emphasis on how well that application will run over the WAN.
- Most network groups have little if any application expertise.
- Until recently, it was rare to find an organization that had ongoing responsibility for application delivery.
- Today, roughly two thirds of network groups have at least some responsibility for application delivery.

Who has Responsibility for Application Delivery


Group Responsible	Percentage of Respondents
Network Group	64.6%
Application Development Group	48.5%
Server Group	45.1%
Storage Group	20.9%
Application Performance Group	18.9%
Other	12.1%
No Group	6.3%

The Factors That Make This Important and Difficult

- In some instances, the IT organization acquires the application from a third party and so cannot influence how it is developed or modified.
- Application delivery involves multiple technologies and organizations; i.e., applications, WAN, LAN, SAN, servers, data bases, security, OSs.
- Other factors that can influence application performance include the performance of the desktop and routing issues such as routing flaps.
- Applications are becoming more complex; i.e., n-tier applications, SOA and Web services.

Web Services Based Applications

- In a Web services based application architecture, the application is composed of multiple modules – referred to as Web services.
- The servers that run the Web services may be housed within:
 - A given data center owned by the enterprise.
 - Multiple data centers owned by the enterprise.
 - Multiple data centers owned by different entities.



The Factors That Make This Important and Difficult

- Consolidating servers out of branch offices and into centralized data centers.
- Data center consolidation.
- Hosting an application in a single data center.
- The increasingly distributed organization.
- The deployment of protocols that are chatty (i.e., CIFS), dense (i.e., XML) or computationally intense (i.e., SSL).

Agenda

- The Application Delivery Track
- The Current Approach to Application Delivery
- Factors that Make Application Delivery Difficult
- A Framework for Application Delivery
- Suggestions and Summary



Application Delivery Framework

- There are four components of Application Delivery. They are:
 - Planning
 - Network and Application Optimization
 - Management
 - Control



Planning

- Planning functions include:
 - Profile an application prior to deploying it.
 - Baseline the performance of the network.
 - Establish goals for the performance of the network and for at least some of the key applications that transit the network.
 - Model the impact of deploying a new application.
 - Identify the impact of a change to the network, the servers, or to an application.
 - Develop a network design that maximizes availability and minimizes latency.

Application Profiling

Network Latency	Trans 1: Goal = 5 seconds	Trans 2: Goal = 5 Seconds
0 ms	2	2
100 ms	2	3
150 ms	2	4
250 ms	2	7
350 ms	4	18
450 ms	4	34
500 ms	12	57

Network Baselineing

- The goal of baselining is to quantify the characteristics (e.g., response time, utilization, delay) of applications and IT resources including servers, WAN links and routers.
- Baselining allows an IT organization to understand the normal behavior of those applications and IT resources.
- Baselining is an example of a task that one can regard as a building block of management functionality. That means baselining is a component of several key processes:
 - Performing a pre-assessment of the network prior to deploying an application
 - Performing proactive alarming
 - Capacity/latency planning

Baselining & Capacity/Latency Planning

- Less than half (43%) of IT organization baseline their network.
- Just less than two-thirds of the companies that do baseline their network (63%) claim that they do it well.
- Roughly 50% of IT organizations perform capacity planning.
- Just less than two-thirds of the companies that perform capacity planning (64%) claim that they do it well.

Application Delivery Framework

- There are four components of Application Delivery.

They are:

- Planning
- Network and Application Optimization
- Management
- Control



Network and Application Optimization

- The primary goals of network and application optimization techniques are to:
 - Reduce the amount of data that is sent over the WAN.
 - Ensure that the WAN link is never idle if there is data to send.
 - Reduce the number of round trips, or application turns that are necessary for a given transaction.
 - Offload computationally intensive tasks from client systems and servers.

Characteristics of the Current Solutions

- There are two general classes of solutions:
 - Branch office focused solutions
 - Data center focused solutions
- The goal of a branch office optimization solution is to improve the performance of applications delivered from the data center to the branch office over the corporate WAN.
- These solutions implement techniques such as compression, caching, protocol acceleration and are proprietary.
- They typically require an appliance in each branch office and another in the data center.

Characteristics of the Current Solutions

- Data center solutions from hardware vendors involve a class of device usually called an Application Front End (AFE).
- AFEs started out as simple server load balancers (SLB).
- AFEs focus on traffic that comes into the data center over the Web.
- AFEs are standards based.
- Current AFEs still function as a SLB, but also offload computationally intensive tasks off the servers; i.e., TCP multiplexing, processing SSL.

Network and Application Optimization

- There is no solution that optimally meets all needs.
- To deploy the appropriate network and application optimization solution, IT organizations need to understand the problem that they are trying to solve.
- To understand the performance gains of any network and application optimization solution, that solution must be tested in an environment that closely reflects the environment in which it will be deployed.

Application Delivery Framework

- There are four components of Application Delivery.

They are:


- Planning
- Network and Application Optimization
- Management
- Control



Management

- Some of the key management tasks are:
 - Discover what applications are running over the network and identify who is using them.
 - Gather the appropriate management data on the performance of the applications and the infrastructure that supports them.
 - Provide end-to-end visibility into the ongoing performance of the applications and infrastructure.
 - Identify the sources of delay in the performance of the applications and infrastructure.
 - Automatically identify issues and resolve them.

Sources of Application Degradation



Existence of Formalized Processes for Identifying & Resolving Application Degradation

Response	Percentage of Respondents
Yes, and we have had those processes for a while	22.4%
Yes, and we have just recently developed these processes	13.3%
No, but we are in the process of developing these processes	31.0%
No	26.2%
Other	7.1%

Impediments to Application Delivery

Impediment	Percentage of Respondents
The processes that we have are inadequate	39.9%
The difficulty in explaining the causes of application degradation and getting any real buy-in	35.1%
The tools that we have are inadequate	32.7%
There is an adversarial relationship between the applications development group and the rest of IT	23.6%

Discovery

- Traditional management adage: You cannot manage what you can not measure.
- Application Delivery Adage: You cannot manage what you cannot see.
- Only 23% of IT organizations claim that they do a good job of identifying the applications that are running on their network.
- Roughly the same percentage claim that they do a good job of identifying non-approved and inappropriate applications such as Internet radio.

End-to-End Visibility

- Enables IT organizations to notice application performance degradation before the end user does.
- Identifies the correct symptoms of the degradation.
- Facilitates making intelligent decisions and getting buy-in from other impacted groups.
- Allows the IT organization to measure the performance of critical applications before, during and after it makes changes.

Setting Static Performance Alarms

Parameter	Percentage
WAN Traffic Utilization	81.5%
Network Response Time (Ping, TCP Connect)	58.5%
LAN Traffic Utilization	47.8%
Application-Response Time (Synthetic Transaction-Based)	30.2%
Application Utilization	12.2%
Other	5.9%

Approach to Setting Thresholds

Approach	Percentage of Companies
We set the thresholds at a high-water mark so that we only see severe problems.	64.3%
We set the thresholds low because we want to know every single abnormality that occurs.	18.3%
Other (Please specify).	17.4%

Application Delivery Framework

- There are four components of Application Delivery. They are:
 - Planning
 - Network and Application Optimization
 - Management
 - Control



Control

- Some of the key control tasks are:
 - Ensure the availability, security and performance of the desktop.
 - Authenticate traffic.
 - Provide call control and signaling functionality.
 - Classify traffic based on myriad criteria.
 - Prioritize traffic that is business critical and delay sensitive.
 - Perform traffic management and dynamically allocate network resources.
 - Affect the routing of traffic through the network.

Control

- Desktop performance degrades with use due to the accumulation of unnecessary files and applications.
- Routing protocols such as OSPF often do not choose paths with the lowest latency.
- QoS is required in those instances in which the network suffers congestion and there are business-critical applications that are delay sensitive.
- Network Access Control (NAC) seeks to determine if a device should gain access to a network and what resources it can access.
- NAC provides a preview of the coming interaction between security and application delivery.

Agenda

- The Application Delivery Track
- The Current Approach to Application Delivery
- Factors that Make Application Delivery Difficult
- A Framework for Application Delivery
- Summary and Suggestions



Summary and Suggestions

- If you work in IT, you either develop applications or you deliver applications.
- In the vast majority of situations, the end user notices application degradation before the IT organization does.
- The goal of application delivery is to develop the ability to minimize the occurrence of application performance issues and to both identify and quickly resolve issues when they occur.
- Application delivery needs to have top-down approach, with a focus on application performance.
- Successful application delivery crosses myriad technologies and organizations.

Summary and Suggestions

- IT Organizations should:
 - Create an architecture function that integrates application development and application delivery.
 - Involve the application delivery function in the application development/selection process – preferably early in the process.
 - Implement tools that have creditability throughout the IT organization.
 - Improve the processes that support application management.
- More information: *The Application Delivery Handbook*, www.kubernan.com

Thank You!!!