



SERVICE DESCRIPTION WAN PERFORMANCE OPTIMIZATION

Visibility – What You Can't See *Can* Hurt You



This document has been prepared by Predictive Networks solely for service description purposes.

INTRODUCTION

This document outlines the service description of Predictive Networks' Application Performance Optimization services.

The objective is to enable corporation to obtain more granular insight into application performance and network load characteristics of the WAN connected locations, through a centralized device deployment.

Today's wider issue of application delivery and visibility has become very complex.



Many different types of applications are demanding resources from the shared Wide Area Network services. It has become increasingly more difficult to distinguish business from recreational traffic. It is becoming apparent that increasing bandwidth capacity is not a solution for seamless delivery of business critical applications over the WAN.

Research by the Aberdeen Group has shown in their WAN Optimization report that issues such as visibility into application performance and end-user experience as well as managing application performance in virtualized environments is becoming more prevalent. It also showed that Application performance Management increases in bandwidth requirements due to rollouts of new applications and lack of visibility into application performance were the top two concerns that organizations have for application performance management.

Predictive Networks has build a unique and state of the art solution addressing Visibility, Control and Optimization of WAN-based application infrastructures. Predictive Networks has the technology, experience, and services in place to manage performance in a structured way and achieve the following goals such as:

- Achieve an improved and more consistent user experience for some business critical applications globally
- A more cost efficient capacity planning of network bandwidth
- Improve on reactive and pro-active fault management regarding application performance and availability
- Increased visibility/understanding of application behavior and performance

Areas of Cost Savings and Cost Mitigation

This document details how we deliver insight into WAN application usage characteristics enabling corporations to

- determine the true operational usage per remote site per application
- review the monthly cost of the WAN
- avoid unnecessary investments in application acceleration technology and associate maintenance cost
- reduce support staff.

Each assignment entails a modular approach; the implementation will include a baseline of the WAN application infrastructure using a single application visibility probe on the WAN connection between the server network and the WAN and professional services with data analysis and a monthly report and can be expanded to multiple deployment of application performance optimizers and accelerators.

Enjoy reading this positioning paper

Kind regards,
Predictive Networks

Wilfried van Haeren
President

1. UNIFIED PERFORMANCE MANAGEMENT

Take Charge of Your Network and Applications

With network traffic growing exponentially, application performance and WAN optimization are two of the biggest IT challenges that organizations face. Until now, network managers and administrators lacked a comprehensive solution for monitoring and managing network resources and bandwidth while meeting user expectations for fast application response. Exinda revolutionizes the way organizations manage their network by offering an affordable and easy to use appliance that delivers visibility, control and optimization.

Unified Performance Management

Many organizations rely on a number of different point solutions for network monitoring, reporting and application acceleration. This approach adds complexity and increases the administrative burden on IT staff.



With the Unified Performance Management (UPM) solution, all of the core capabilities needed to effectively manage a WAN are incorporated into a single network appliance. These tightly integrated capabilities include real-time monitoring, reporting, traffic control, optimization and intelligent acceleration. Users can toggle between a simplified or advanced interface and may also choose a centralized management platform for controlling multiple appliances across the network.

The outcome of UPM is to enable corporations to obtain more granular insight into application performance and network load characteristics of the WAN connected locations, through a centralized device deployment.



Predictive Networks as Gold Partner of Exinda has the technology, experience, and services in place to manage performance in a structured way which results into:

- An improved and more consistent user experience for some business critical applications
- A more cost efficient capacity planning of network bandwidth
- Increased visibility and understanding of application behavior and performance.

2. OVERVIEW EXINDA

Exinda is a global provider of WAN optimization and application acceleration products. The company has helped over 2,000 organizations worldwide reduce network operating costs and ensure consistent application performance over the WAN.

With a single Exinda appliance, organizations can effectively monitor, manage and optimize their entire network. Unlike other solutions, Exinda requires no additional modules or add-on components. This makes Exinda not only more affordable but faster to implement and easier to manage. The complete, purpose-built Exinda solution allows network managers to spend their time optimizing the network instead of managing the various point solutions themselves.

The Exinda Value Proposition

The proposition consists of

- Integrated network visibility, control and optimization in a single “plug-and-play” appliance
- Affordable, all-in-one solution accelerates deployment, reduces complexity and lowers total cost of ownership (TCO).
- Market-leading support for Microsoft Active Directory ensures fine-grained visibility and control down to the individual user level
- Service Delivery Point (SDP) provides centralized management from a single interface

The hardware platform consists of the x700 series for monitoring and optimization and the x800 series, which includes Acceleration.

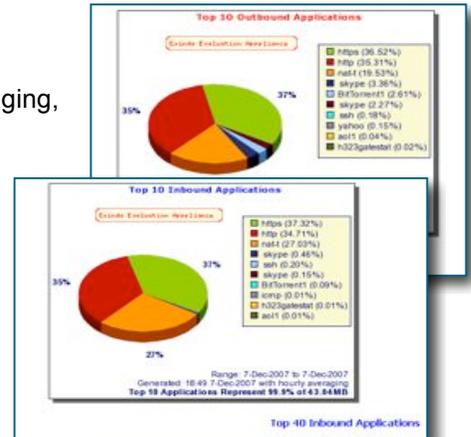


On October 12, 2009 Exinda announced its new top-of-the-line WAN optimization platform, the Exinda 8760. It tops out at 5 Gbps, the first WAN Optimization devices to break the 1 Gbps barrier. With Exinda’s architecture, customers can scale from 100 Mbps to 5 Gbps on a single appliance that supports a range of optimization features, including application visibility, QoS, protocol optimization and acceleration.

The Unified Performance Management program by Exinda consists of:

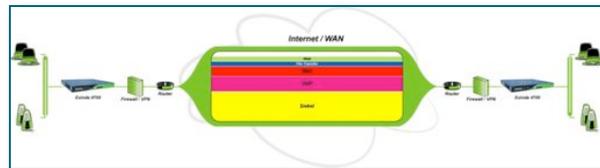
- **Application Visibility**

- Layer 2 through 7 Monitoring
- L7 Signatures – P2P, HTTP, VoIP, Citrix, Messaging, Streaming, more
- Behavioral Based Classification - Encrypted Apps – Bit Torrent, Skype
- Deep Drill Down Inspection – correlate hosts/application usage and breakdown
- Top Talkers, Conversations
- Real Time Monitoring
- Application Response Time Measurements
- Application Specific Analysis Modules



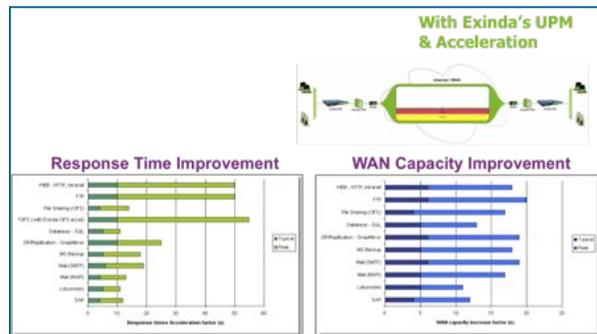
- **WAN Optimization supporting**

- Layer 2 through 7 classification
- Precise Policy-based Traffic Management and Shaping (QoS)
- Diffserv Packet Classification and Marking
- Hostname, URL, HTTP file Wildcard classification
- Traffic Discard/Block
- Time of Day Policies
- Adaptive Response



- **Application Acceleration supporting**

- Wide Area File Services (WAFS)
- TCP Acceleration
- CIFS Acceleration
- WAN Memory
- Compression



- A Scalable **Centralized Management** environment using SaaS based portal technology.

- Hosted service provided by Predictive Networks based out of London, ON
- Provision, configure, and manage devices
- Centralized reporting from multiple Exinda appliances
- Dashboard views to consolidate WAN wide data
- Immediate or scheduled configuration changes and software updates
- Multi user support with role based administration

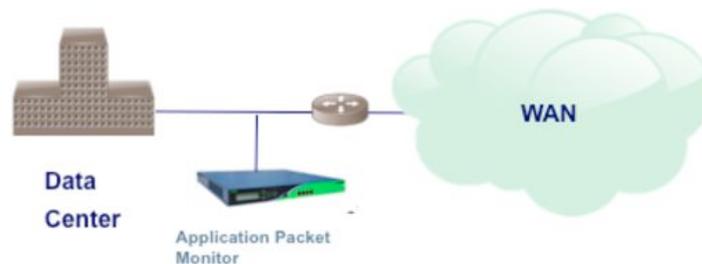
3. EXINDA DEPLOYMENT

Visibility starts with monitoring the current application flows across a Wide Area Network. Devices very suitable for this task are the EX x700 series typically installed at a central data center site.

Deployment outline

The hardware devices available range from 2 Mbps link devices to 5 Gbps high-end monitoring systems.

In monitoring-only mode, it will be deployed as a non-intrusive configuration

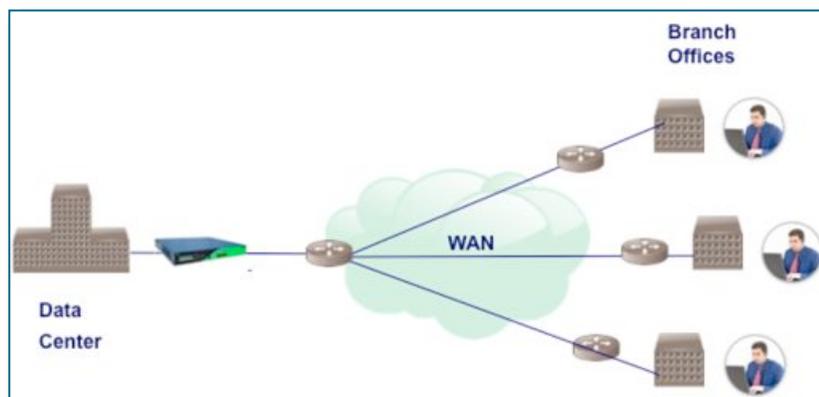


This topology is used when customers need to monitor only, without installing the Exinda in in-line mode. The Exinda will monitor and report on all applications presented on the SPAN/mirror port. This is regularly used to perform network audits as it provides great flexibility in restricted and complex network environments.

FROM MONITORING TO OPTIMIZING

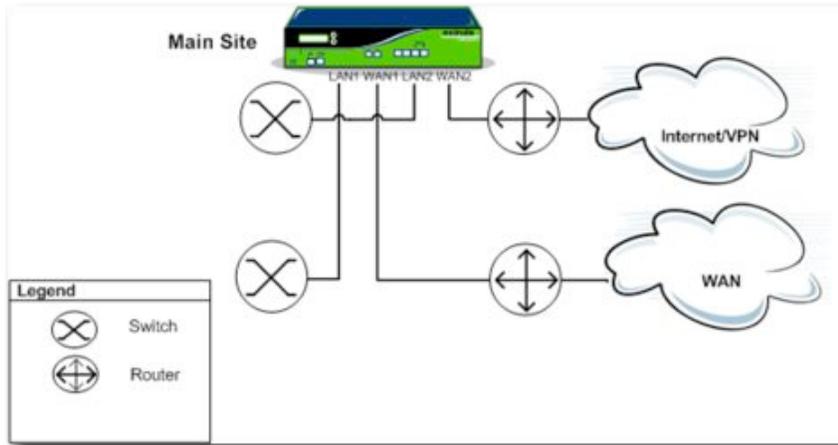
An x700 series appliance as monitoring-only device can be upgraded to an application optimizer by positioning the appliance between the central switch and the outbound router enabling the appliance to work in **optimization mode**. This allows the activation of QoS on the basis of policies and application signatures. Quality of Service is an extremely powerful technique that maximizes the performance demanded by users and applications.

In optimization-mode the appliance will be deployed in-line. In many infrastructures the WAN topology is of a 'hub-and-spoke' model whereby all data flows from remote site to central data center. A single device deployed at the central location is sufficient to monitor and control all flows.



Web-based corporate applications, Thin Client and Telnet sessions can take advantage of a higher priority than a less time critical service such as File Transfer Protocol (FTP) or Email. Corporations may wish to allocate amounts of bandwidth to different divisions or business units such as a guaranteed amount of bandwidth to customer service or manufacturing applications and users. Often this reduces costs by eliminating networks designed for peaks in traffic rather than averages whilst simultaneously decreasing response time for users.

The high-end x700 appliances support multiple WAN interfaces, offering the monitoring and control of private WAN and Internet traffic to and from the main site.



FROM MONITORING & OPTIMIZATION TO ACCELERATION

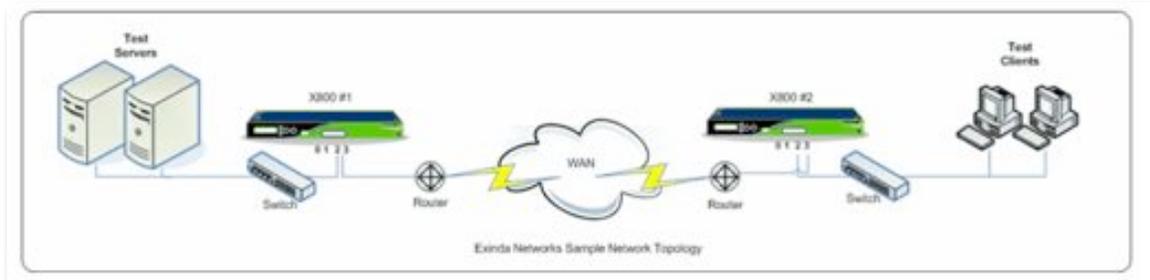
In acceleration deployment, requiring 2 or more Exinda devices, the appliances will use traffic traversing the WAN to auto discover each other. Following auto discovery, and successful version negotiation, they will begin to accelerate applications between them. The acceleration happens based on Exinda's default policies. These policies can be altered using the management portal. The x800 series enable next to monitoring and optimization, the acceleration functionality.

To Accelerate or not to Accelerate

Some applications simply don't benefit from application acceleration, but do benefit from QoS. Secure applications like ssh or VoIP applications for example will not benefit from acceleration. For these applications, Exinda appliances will prioritize these and make sure they get the service they deserve.

Typical protocols that Exinda accelerates are: CIFS, FTP, HTTP, Data backups (unencrypted), Lotus Notes, Novel, Some applications that Exinda only applies QoS to are: VoIP, Citrix and other thin client applications

The appliances include a combination of QoS and acceleration policy templates. In a QoS + Acceleration template Exinda will apply a best practice policy set for controlling and accelerating applications.



The **x800 series** of Exinda offer next to monitoring and optimization also acceleration.

4. SERVICE DELIVERY POINT

Exinda's Service Delivery Point (SDP) is available through Predictive Networks as hosted "Software as a Service" and licensed by the number of appliances to be managed. The hosted service further simplifies the WAN optimization process by eliminating the need to buy, install, maintain and upgrade management hardware and software. Since there is no capital investment, SDP offers a 70 - 90 percent savings compared to similar client/server solutions.

SDP outshines traditional centralized management systems (CMS) by eliminating the cost, risk and complexity of such implementations. SDP provides a quicker return on investment as it can be deployed quickly, typically in one week versus three months or more for a traditional CMS. SDP provides full visibility into network usage at application layer 7 while most CMS only focus on network layer 3.

Predictive Networks as **Gold Partner** offers the SDP appliance hosted at it's secure data center facility in London, ON, Canada.

Exinda's Closed Loop Networking (CLN), a combination of Exinda's traffic management appliances and Exinda's SDP, assists CIO's with addressing their three highest priority business initiatives:

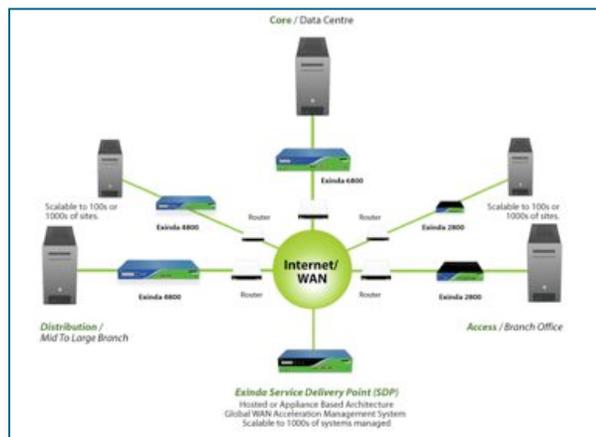
- Increasing user productivity due to consistent service levels
- Providing service continuity
- Reducing ongoing operational costs



CLN and SDP

Networks are under constant traffic congestion from internal business applications, internal non-business applications and unwanted external applications. Network managers struggle to keep pace with this ever increasing number of application flows and, therefore, find it difficult and resource intensive to maintain appropriate service levels.

CLN and the functions that SDP offers within the CLN architecture provide an automated process to address this predicament. Continual monitoring, auto-adaptive responses and steering adjustments provide computer controlled, comprehensive management of networks. The software dashboard and reporting function is available through the Service Delivery Point. In the pilot only one device will be deployed. The capture data is pushed to the SDP using secure HTTP (HTTPS) file transfer.



5. APPLICATION RESPONSE-TIME MEASUREMENT

Traditional latency measurement technologies provide a one-dimensional view of the actual cause of the problem. A typical example of such a technology is ICMP echo request/response, more commonly known as *ping*. Ideally, ping is a computer network tool used to test whether a particular host is reachable across an IP network. However, using interval timing and response rate, ping also estimates the round-trip time, which is essentially a measure of network latency.

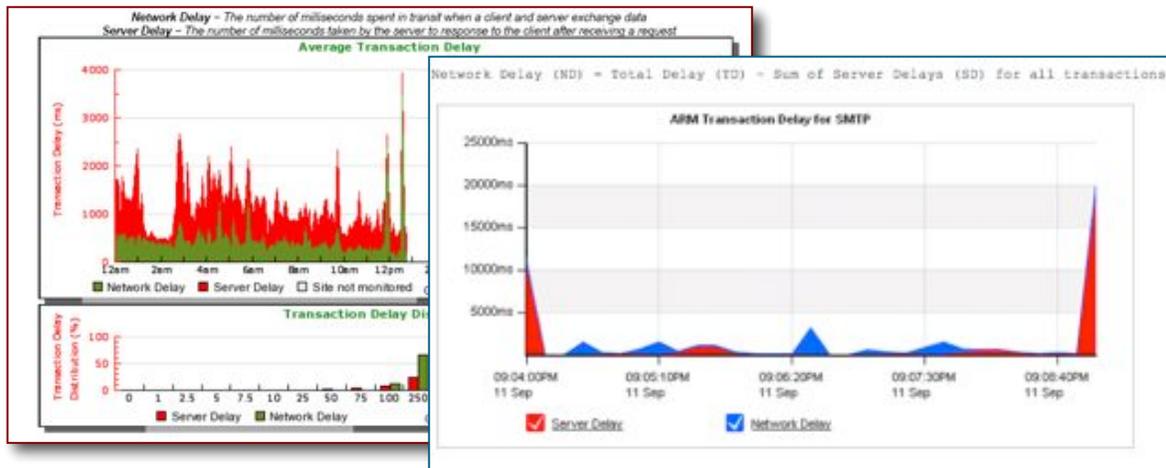
Another technology that might help is monitoring the server load using Simple Network Management Protocol (SNMP) server monitoring. Monitoring server load during times of poor application response might reveal any server load related performance issues.

These technologies however, only provide an insight into the problem of slow application response without identifying the real cause of the problem.

Application Response Measurement

Application Response Measurement (ARM) is a tool that helps get to the root of the problem, identifying the cause of the delays. ARM passively monitors TCP transactions to get estimates of delays on the network and server. The results of the monitoring are broken down and presented in a way that the cause of the delays is obvious to the administrator. ARM enables organizations to

- Quantify application performance from the end user perspective
- Detect how long end users are waiting for their applications to respond
- Pinpoint if a problem is network or server related
- Help fine-tune QoS policies to control response times.



ARM enables application support teams to view their application behavior, characteristics and pay load from an infrastructure angle. Using SDP and **role-based logins**, apps teams can be offered role specific dashboard information.