

Performance,
Scalability,
&
Integration

Valencia Systems Aruba Suite™ for Service Level Management

Valencia Systems
1040 Woodcock Road,
Suite 100
Orlando, FL 32803

Phone: (1) 407.228.4417
Fax: (1) 407.228.4419

www.valenciasystems.com

Table of Contents

Section 1: Introduction

Section 2: Components of Service Level Management

Section 3: Aruba Suite Overview

Section 4: Service Level Management (SLM) Reports in Action

Section 5: Summary

Section 1: Introduction

IT organizations are tasked to maintain uptime and limit the downtime for network services. When problems are detected the average time it takes to fix the problem, mean-time-to-repair (MTTR), needs to be minimized. Each second of downtime corresponds to lost worker productivity and real business cost. A common way to ensure that performance is met is to set a predefined measure of what the minimum level of performance needs to be to allow users to support their network requirements. These Service Levels can be formalized agreements with internal or external providers or IT guidelines which need to be managed and maintained.

Section 2: Components of Service Level Management

Key areas of relevance for supporting Service Levels are:

- Identifying who is providing the service, internal or external
- Existing Service Benchmarking
- Identifying key service components
- Setting the service level, availability and performance
- Monitoring and managing the Service
- Troubleshooting
- Reporting

Who is responsible?

Before any management system can be in place, the responsible organization or organizations needs be identified. The service may be provided by an internal department, i.e. IT, or from an outsourced provider, or a combination of both. Depending on the applications provided the scope of management can become quite complicated, very quickly. An example could be an email service. In a simple environment the IT department has responsibility and ownership for the service delivery including the email server, the network and the application residing on the desktop. When email become unavailable it is easy to identify that there is an IT service problem and that the minimum service level is not being met. Even the troubleshooting of the service is fairly easy when you have access and control to all of the parts. The task becomes more complex as parts of the service are outsourced to managed service providers of applications and network service. Setting up service levels and determining breakdowns requires holistic visibility into the underlying infrastructure including the network devices, systems and applications.

Benchmarking the service?

To determine how to establish a service level it must be determine what an acceptable service level is. Does the service need to be available 99999% of the time? What is the performance of the application to the actual user? Are there financial agreements negotiated based on the service? These are all considerations that need to be addressed when evaluating and managing a service level.

Identifying key service components

It is key to identify the vital service components that make up a service. What systems are supporting the service delivery? What network components are involved? How is the network currently being utilized? What other application traffic is using the network? Are there any security issues? What is the IT roadmap for capacity planning? What business units and locations are using the resources? Who is using the service?

Setting the Service Level

The service level can be set based on a contracted level, i.e. the network needs to be available 98% of the time, exempt from planned outages. The service level could be more sophisticated in managing a predefined service or application, i.e. Email needs to be available 99% of the time, except for planned outages. The service level can become even more specific by insuring that Email needs to be available 9x% of time with performance less than xMS. Priorities may also be set to insure that individuals or departments have priority of level service for specific applications or general availability.

Monitoring and Managing the service

Service Levels can become quite complicated, rather quickly. The framework to support Service Level Management starts with a scaleable reporting architecture and the ability to monitor remote service performance. The underlying systems required to assure that service levels are met require the capacity to collect, analyze, store, present, and augment vast amounts of consolidated material. Most solutions, today, portray their solutions as scalable but they lack the internal architecture needed to grow. Most vendors have concentrated on attacking the problem in bite size pieces and presenting segmented attributes that focus on individual aspects of service quality. The problem needs to be viewed in a holistic manner, rather than analyzing the parts. A solution should allow the management of complex service levels, in complex (distributed) network environments, trending, historicists, and integration.

Trouble Shooting-Problem analysis

Once a service level is exceeded the underlying problem needs to be quickly identified so resolution can take place. Downtime is not acceptable in a corporate environment and service level management solutions need to provide quick analysis on not just that there is a failure, but where and why the failure has occurred.

Reporting: Historical, Segmented, User-Based, Business Intelligence

The historical analysis of service levels for tracking the Quality of Service (QoS) allows the organization to “score card” how IT is performing on a daily, weekly, monthly and annual basis. The ability to roll the daily information up and track the ongoing QoS provides a management index on the overall health of service performance. This specific analysis provides management with the ability to make top-level decisions in capital budgeting, outsourcing, and future company IT requirements. User based reporting allows IT to segment the reports for individual users, user groups or as a managed service view. These groups can be branch, division, customer based, executive level or IT based; customized to the specific user group for clear and concise presentation.

Business Intelligence

Business Intelligence (BI) is the ability to gather data from multiple systems in raw form and turn it into intelligence that business decisions can be formulated. It is the umbrella that ties together other closely related data disciplines including data mining, statistical analysis, forecasting, and decision support. Business Intelligence is often the top layer of collected information, which allows business managers to make IT capital expense decisions or provides information on how to optimize existing infrastructure. The presentation of Business Intelligence requires flexible and customizable reporting, since the information will be shared and presented to potentially many lines of business and levels within an organization

Section 3: Aruba Suite Architecture & Features



The Aruba Suite leverages a Highly Scalable data collection capability and database to provide a "carrier-class" consolidated IT performance management reporting solution. Aruba collects extensive raw SNMP data from hundreds of different networked devices to isolate "hot spots" and report on real-time conditions for troubleshooting, Capacity Planning, Service Level Management (SLM), Business Service

Management (BMS), Quality of Service (QoS), financial-bill back, and operational needs. Raw data is rolled-up into historical reports for current time periods or multi-year analysis. Aruba is easily configured and installed. Aruba ships with hundreds of preconfigured reports and can be easily customized (using XML) for individual requirements for specific users (IT Managers, Operations, Engineering, etc). Aruba's powerful base lining capability allows enterprises and service providers to determine typical vs. atypical behavior, trending and predicting behavior and identifying security or policy breakdowns. Aruba is designed to be integrated with existing IT applications to provide a comprehensive, scalable, IT performance management solution.

Aruba Product Features

The Aruba Report/Server

- Carrier-class scalability
- Historical trending & predictive analysis
- Hardware & network capacity planning
- Comprehensive System Management reporting
- Resource utilization & optimization reports
- Efficient provisioning of carrier capacity reports
- Proactive problem area detection
- Verify/Audit QoS goals
- Verify/Audit SLA goals
- Automatic Baseline comparisons
- Threshold based exceptions and alarms
- Secure Access
- ODBC compliant database
- XML- customization

The Aruba Distributed Flow Collector

- Unprecedented scalability, millions of concurrent flows

- Automatic Baseline comparisons
- Service Level (SLA) assurance
- Measure service quality (QoS)
- Utilization
- Congestion
- Errors
- Real-Time Analysis
- Historical trending & predictive analysis
- XML-Web based reporting
- Comprehensive picture of Network Usage, raw & aggregate data
 - Top N Talkers, Hosts, Listeners
 - Top N Conversations
 - Top N Ports
 - Top N Applications

The Aruba Distributed Poller

- Distributed remotely or centralized
- Add modularly as requirements grow n-tier architecture
- Multithreaded, high speed, polling of SMNP data cache
- Standards-based SNMP polling

The general Aruba Architecture is a scalable architecture which supports the remote/distributed collection of data from routers, hubs, switches, computer systems and other SNMP capable devices (i.e. firewall, or traffic shaping appliances). By distributing the polling function the data may be consolidated from the distributed collectors in a central data warehouse. Data is analyzed against preconfigured or customized report formats to produce actionable, information-based, web reports. Reports are easily categorized and grouped by function, location, priority and detail allowing Network Operations-to-Management utilize performance information for IT troubleshooting or financial and business planning.

Aruba was specifically architected to support very large-scale SNMP/NetFlow data collection and reporting. The architecture is both multi-tiered and distributed. Aruba can be divided into three logically separate layers: **Client layer**, **Server layer**, and **Collection layer**.

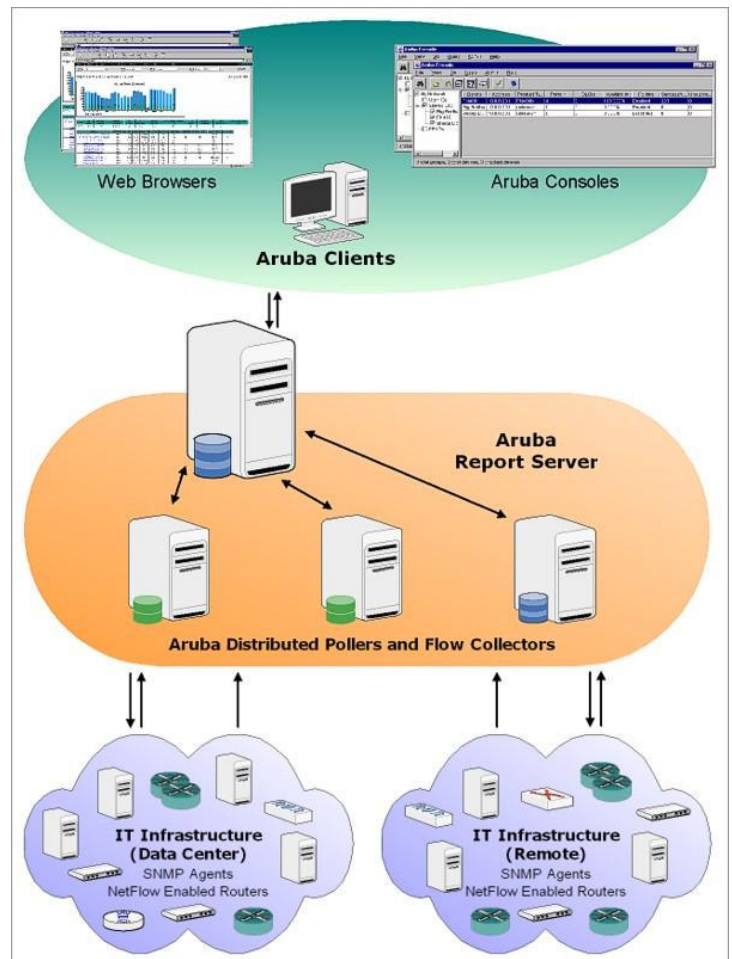
Aruba is a highly scaleable architecture utilizing a client, server, collection layers and an Oracle™ database

The Client layer provides the user interface to the system. The Aruba Console allows the administrator to configure and maintain the system, as well as perform real-time and custom reports.

The Server layer (*Aruba Report Server*) is responsible for generating and publishing reports, storing configuration information, and managing the underlying Pollers and Flow Collectors. There is a single Report Server in an Aruba system.

The Collection layer consists of one or more Pollers and/or Flow Collectors which automatically gather data from the network devices. These clean, normalize and store raw data, as well as aggregating and aging historical data.

The Client, Server and Collection may all be collapsed onto a single machine or they may be distributed on separate computers. For small networks, a single Client, Report



Server, and Poller running on the same machine will be appropriate. Larger networks may require multiple Clients simultaneously connecting to the Report Server, that in turn, managing data from several distributed Pollers and/or Flow Collectors. Aruba's architecture will scale from small to very large distributed installations, without sacrificing simplicity and ease of use.

Why a multi-tiered, distributed architecture?

The optimal solution must minimize data movement by pre-processing the data at the collectors

A monolithic approach, a data collection engine and reporting engine both hosted on a single platform, will eventually be limited by processor performance, available bandwidth, and storage capacity. And long before these thresholds are reached, the user experience is impacted - continuous polling and database processing will consume much of the available machine cycles. Just adding more machines introduces a new problem - each machine now has its own collection and reporting domain. The user must know which machine is responsible for each device, and groups of devices can not span the domains.

An alternative approach is to separate the administrative and reporting functions from the data collection engines. This still gives the user a single machine to access for a holistic view of the entire network, for both reporting and administration. The CPU/bandwidth-intensive collection functionality is distributed across as many machines as needed. An additional benefit is that a collector can now be located close to the data sources, reducing SNMP polling across the WAN.

The downside of this approach, if the collectors rely on the server layer to actually process and store all the data that is collected, is the problem has not been solved. In fact, this only makes it worse, since the data is essentially moved twice. Many solutions which started with a monolithic architecture have been "scalability enhanced" by adding a "server layer" using this approach and must manipulate the data across platforms.

Valencia System feels the optimal solution must minimize data movement by pre-processing the data at the collectors - only summarized data is sent up to the server layer. Granular data (for example, raw polled data) is retained in the collector database, and is only moved to the server if user requested.

Reducing the ongoing maintenance burden is a major factor in a solution's scalability. To reduce maintenance burden the client layer must be distributed as well. The administrators console should tear away, and multiple administrators must be able to simultaneously make changes. In addition, grouping must be integrated into the administrative functions as well as the reports.

Section 4: SLM Reports in Action

Valencia Systems offers a broad array of fully customizable report templates. The objectives of the SLM reporting is to identify who is providing the service (internal and external), benchmarking and measuring performance, setting service levels, troubleshooting and trending (historical analysis) over time.

Examples of SLM reporting are included below:

Wide Area Network SLA Summary

Daily WAN Service Quality Report for North East

Tuesday, April 12, 2005

Collection Hours

Group	Average Availability	Average Daily Uptime	Average Daily Downtime	Data Unavailable	Server Down	Interface Failures	MTTR	MTBF	Average Round Trip Delay	Average Daily Interface Count
North East	100.000%	18h 30m	0	0	5h 30m	0	No Failures	No Failures	177 ms	20


Name	Address	Interface Type	Speed	Availability	Daily Uptime	Daily Downtime	Data Unavailable	Server Down	Interface Failures	MTTR	MTBF	Average Round Trip Delay
157.130.4.58_1	157.130.4.58	LAP-B	1.5M	100.000%	18h 30m	0	0	5h 30m	0	No Failures	No Failures	189 ms
157.130.4.58_4	157.130.4.58	Proprietary Point-To-Point Serial	0	100.000%	18h 30m	0	0	5h 30m	0	No Failures	No Failures	154 ms
157.130.6.114_3	157.130.6.114	Frame Relay	320K	100.000%	18h 30m	0	0	5h 30m	0	No Failures	No Failures	181 ms
157.130.6.114_4	157.130.6.114	Frame Relay	320K	100.000%	18h 30m	0	0	5h 30m	0	No Failures	No Failures	203 ms
CS-New_Jersey_Warehouse-USA_To UUNET (u89425	157.130.2.170	Frame Relay	1.5M	100.000%	18h 30m	0	0	5h 30m	0	No Failures	No Failures	152 ms
Fel-Lodi-INET.IT.felician.edu_To UUNET (u52192)	157.130.10.134	Frame Relay	1.5M	100.000%	18h 30m	0	0	5h 30m	0	No Failures	No Failures	156 ms
GW.HDPublishing.com_MFR1	157.130.6.214	Frame Relay	3.1M	100.000%	18h 30m	0	0	5h 30m	0	No Failures	No Failures	96 ms
Mahwah Internet_2	157.130.10.70	Frame Relay	1.5M	100.000%	18h 30m	0	0	5h 30m	0	No Failures	No Failures	171 ms
Mahwah Internet_3	157.130.10.70	Frame Relay	1.5M	100.000%	18h 30m	0	0	5h 30m	0	No Failures	No Failures	162 ms
Mahwah Internet_4	157.130.10.74	Frame Relay	1.5M	100.000%	18h 30m	0	0	5h 30m	0	No Failures	No Failures	155 ms
Mahwah Internet_7	157.130.10.70	Frame Relay	1.5M	100.000%	18h 30m	0	0	5h 30m	0	No Failures	No Failures	119 ms
Super Duper Parrot Router_1	157.130.12.86	Frame Relay	64.1K	100.000%	18h 30m	0	0	5h 30m	0	No Failures	No Failures	293 ms
the-dms-gw.ALTER.NET_To UUNET (u09585)	157.130.12.18	Frame Relay	1.5M	100.000%	18h 30m	0	0	5h 30m	0	No Failures	No Failures	187 ms
UUNET-3640a_LINK TO UUNET	157.130.8.158	Frame Relay	1.5M	100.000%	18h 30m	0	0	5h 30m	0	No Failures	No Failures	253 ms

Local Area Network Performance Report

Daily Ethernet Latency History Report for the-dma-gw.ALTER.NET_To Pix Firewall **Friday, April 8, 2005**
Collection Hours


This report will automatically refresh after 5 minutes.

Address	Interface Type	Speed	Vendor	Description
208.193.72.49	ethernet-csmacd	100M	cisco	(6) FastEthernet0/0



LAN Utilization Report by location

Daily LAN Utilization Distribution Report for Top Groups **Monday, April 11, 2005**
Collection Hours



ID	Group	ID	Group	ID	Group
1	Florida	4	Central	7	Mountain States
2	Washington DC	5	All Locations	8	South East
3	North East	6	Servers	9	California

Troubleshooting--Anomalies Report-- Major changes to norm

Daily LAN Top 10 Changes from Baseline Report for All Locations Friday, April 8, 2005 Collection Hours

Top 10 Increase over Baseline Volume

Component	Rx Octets	Tx Octets	Total Octets	Baseline Total Octets	Baseline Minimum Total Octets	Baseline Maximum Total Octets	Net Change From Baseline	% Change From Total Baseline
Holiday Inn-LasVegas LAN	185.5M	3.8G	4.0G	394.2M	215.4M	474.4M	3.6G	920.93%
Rmon194_4	170.4M	16.9M	187.3M	60.3M	45.9M	75.3M	127.0M	210.57%
Flowpoint_IDSL_STATIC_1	4.1M	67.4M	71.5M	23.4M	1.1M	38.6M	48.1M	205.95%
Cisco1720_connected_to EthernetLAN	47.0M	373.7M	420.7M	138.0M	66.0M	224.2M	282.7M	204.77%
Rmon194_5	599.5M	394.5M	994.0M	336.1M	43.1M	1.4G	657.8M	195.70%
Darth Maul_2	2.0G	356.5M	2.4G	1.1G	71.2M	2.5G	1.3G	119.59%
63.91.169.33_1	713.9M	439.2M	1.2G	542.3M	322.2M	963.3M	610.8M	112.63%
63.93.110.77_1	19.0M	133.7M	152.6M	73.5M	58.0M	97.2M	79.1M	107.54%
63.93.110.32_1	95.9M	1.5G	1.6G	762.8M	437.0M	1.3G	810.8M	106.30%
63.93.62.47_1	15.9M	115.1M	131.0M	64.5M	15.8M	121.3M	66.5M	103.02%

Top 10 Decrease under Baseline Volume

Component	Rx Octets	Tx Octets	Total Octets	Baseline Total Octets	Net Change From Baseline	% Change From Total Baseline
PZADS08_3	0	0	0	535.3M	-535.3M	-100.00%
63.93.105.102_1	0	0	0	1.3M	-1.3M	-100.00%
host.ohnodesign.com_3	0	0	0	743.6K	-743.6K	-100.00%
Rmon194_13	2.6M	8.0M	10.6M	21.8G	-21.7G	-99.95%
Rmon194_6	363.8K	5.9M	6.3M	10.2G	-10.2G	-99.94%
dana-gw.ALTER.NET - Lisle 630-271-4860	24.5M	26.9M	51.5M	2.1G	-2.0G	-97.54%
myelos.btgc.com_connected_to EthernetLAN	0	1.2M	1.2M	39.1M	-38.0M	-97.01%
63.93.105.102_2	175.0K	112.5K	287.5K	1.5M	-1.3M	-81.35%
David-Ellis-GW_connected_to EthernetLAN	35.0M	131.5M	166.5M	842.6M	-676.2M	-80.24%
63.93.105.98_2	52.8M	31.7M	84.5M	422.1M	-337.6M	-79.99%

Section 5: Summary

The ability to administer and manage service level agreements can become a complicated problem very quickly. The underlying solutions should offer scalability and a holistic approach to solving the SLM problem. The visibility into the internal or external service infrastructure is a critical requirement in identifying and managing services and equally critical is the need to present that information in a format that can be quickly and easily understood by the target audience. Business decisions on corporate cost and expense and the performance of the underlying IT infrastructure need to be analyzed to make service level management successful.

To learn more about Valencia software solutions, please visit us on our web site at www.valenciasystems.com