Network Management: Implementing and Operating High Availability Solutions

Executive Summary

Preface

This white paper is for network management engineers involved in designing, implementing, and operating highly available network management solutions. It is helpful to have familiarity with Cisco network management products, network management protocols such as Simple Network Management Protocol (SNMP), syslog, Network Time Protocol (NTP), and server hardware and operating systems.

Introduction

This white paper introduces the reader to principles in designing, implementing, and operating highly available network management solutions.

Networks are quickly growing in size as new IP-capable devices, such as wireless access-points and IP phones, proliferate in an environment. Additionally, more companies depend on their networks for their critical business needs. Fault-tolerant network management is needed to help ensure business continuity.

Solution/Technology

This paper first covers high availability network management principals in general terms. Then it covers high-level guidelines for several Cisco network management products. This paper can be used for early solutions planning.

Identify Network Management Functions Needing High Availability

The first step in operating a fault-tolerant and highly available network management environment is to assess which network management functions are in place and which truly need to be highly available.

Network management tools can generally be broken down into the fault, performance, configuration/inventory, availibility, accounting, and security functions :

- Fault management tools: Receive SNMP traps and syslog event messages; some do active polling. These tools are generally prime candidates for high availability. Most devices can handle forwarding SNMP traps and syslog event messages to multiple event receivers. It is generally not advisable to configure more than six SNMP trap and six syslog event message receivers. If additional receivers are needed to feed other network management or security management tools, please review the following sections on redundancy with SNMP traps and syslog event messages using repeaters/multiplexers.
- Performance management tools: Transmit and receive SNMP polling for near-term fault monitoring and long-term performance trending. If this is a high availability requirement, then having multiple applications polling a device for normal device health checks is typically fine. Care should be given that the applications are not doing extensive SNMP walks on low-end devices for objects such as routing tables, ARP caches, MAC address tables, ATA flash disks, and so on. If performance management tools are only used for long-term trending, it may be sufficient to operate them without high availability in mind. If the system goes down, administrators simply lose some time period of trended data.

- Configuration/Inventory management tools: Transmit and receive configuration changes of managed devices. These tools are generally prime candidates for high availability. Multiple Network Management Systems polling the same device for configuration and inventory is generally not an issue, but some devices don't allow multiple, simultaneous operations to Flash memory or simultaneous configuration transfers. To mitigate this, note the options provided in the CiscoWorks LAN Management Solutions (LMS) section later in this document.
- Availability monitoring tools: Transmit and receive up/down polling information, many times emulating mission-critical traffic to gauge availability and latency. These tools are prime candidates for high availability. They must always be more highly available than the environment they are monitoring. Redundant solutions provide an opportunity for maintenance and fault tolerance in the server hardware or software application.
- Accounting tools: Receive accounting/usage statistics; sometimes this is rolled into performance management. Some environments don't do accounting for things such as internal department chargeback, so that functionality would be deprioritized or descoped altogether. However, some environments need accounting information to track for usage statistics related to billing. Providing redundant accounting solutions also requires an ability to manage duplicate entries.
- Security management tools: Transmit and receive authentication, authorization, and accounting (AAA) information and may perform vulnerability tracking services. Losing login authentication and authorization functions is typically traumatic to business continuity so these services are typically mandated to be highly available. Authorization to specific applications, processes, and workflows may be deferred during an AAA outage or permitted by policy if the user has previously authenticated. Accounting of the network device changes being made may be critical to regulatory compliance, so this function may warrant high availability consideration.

Any customer service-level agreements (SLAs) or regulatory compliance guidelines or mandates will likely guide the prioritization of these components.

Redundancy with SNMP Trap Receiving

In a large network management environment there may be multiple SNMP trap receiving applications. A high availability network management environment may extend the need for additional SNMP trap receivers to cover the redundant servers. Common solutions for dealing with multiple trap receivers are:

- Increasing the number of SNMP trap receiver definitions in the managed device configuration
- Using SNMP trap forwarding on one application to receive the alert, allowing it to forward on to another application
- Using an SNMP trap repeater or multiplexer up-stream of the network management applications

The first solution is fine, but it is not recommended to go beyond six SNMP trap receivers in a managed device configuration (Figure 1).

Figure 1. General Fault Management Framework

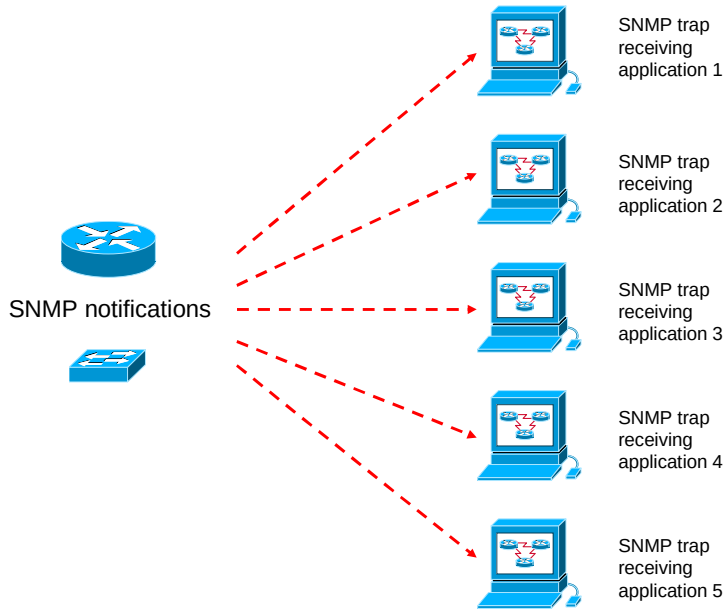

If more than six are needed, it is recommended to look at the other two solutions to limit the amount of processing overhead on the managed device related to transmitting the same alert over and over.

The second solution is acceptable, but consider the scenario where another trap receiving application resides on the same server; it will be required to internally forward the event on a nonstandard port other than User Datagram Protocol (UDP)/162 (Figure 2).

Figure 2. Device to Server and Application Flow

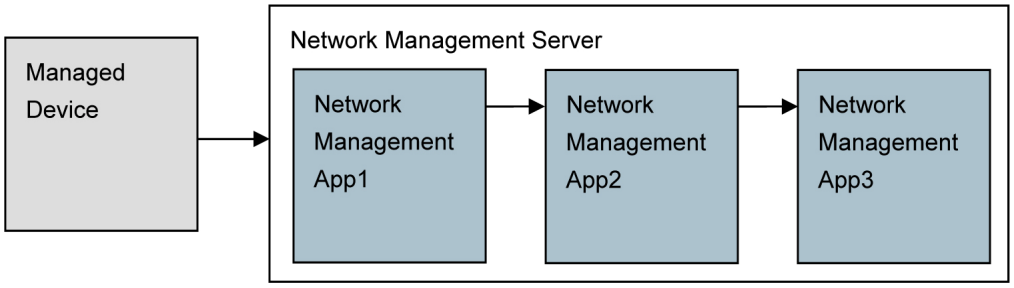

The Managed device would send the initial notification to the network management server where network management App1 initially accepts it as a standard SNMP trap UDP/162.

Network management App1 would forward to network management App2 as an SNMP trap on a nonstandard port, UDP/10162.

Network management App2 would forward to network management App3 as an SNMP trap on a nonstandard port, UDP/20162.

If network management App2 has a failure or is down for maintenance, then network management App3 would lose SNMP trap forwarding and visibility into network faults.

It is not advisable to “chain” too many applications, or the risk of a server or application outage reveals the single-point-of-failure inherent to this solution.

The third solution of using a trap repeater or multiplexer is good if dual SNMP trap repeaters/multiplexers are deployed and monitored (Figure 3).

Figure 3. Fault Management Framework with SNMP Trap Repeater/Multiplexer

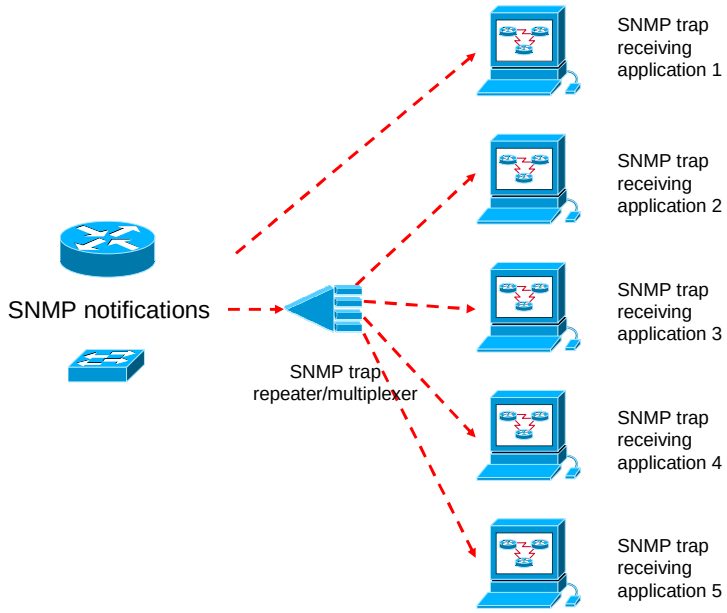

Some SNMP trap repeater/multiplexer options are:

- Commercial
  - Concord eHealth TrapEXPLODER
  - IBM-Tivoli Netcool/Precision IP – TrapMux
  - RealOps TrapBlaster
- Open source
  - BDSI straps (simple trap multiplexer)
  - LooperNG – Event Routing System

Redundancy with Syslog Event Messaging

Syslog event messaging has similar areas of concern as seen with SNMP trap receiving. A management framework design should make sure that “chaining” too many receivers together does not provide unreasonable points of failure.

Many Unix-based operating systems support filtering and forwarding of syslog event messages in their syslog daemon configuration files (syslog.conf), however it should be noted that many syslog daemons will rewrite the header of the message with the syslog server as the message originator as it forwards on to the next syslog receiver. Clearly this is not the desired effect; the originator’s headers must be preserved.

An excellent solution for forwarding/repeating syslog event messages that will preserve headers is the Balabit Syslog-NG product[1]. The initial syslog receiver that forwards/multiplexes the messages and the end-receiving syslog event receiver must be Syslog-NG or the final syslog event receiver must be RFC 3164/3195 compliant.

Identify a Network Management Application’s True High Availability Capabilities

Once the required network management functions that need to be highly available have been identified, the next step is to assess the true high availability capabilities of the network management tools. Many tools have no inherent high availability capabilities. Others may have broad high availability capabilities with redudant application server and/or database server options. It may be possible to take tools that have no inherent high availability capabilities and operate them as isolated pairs to achieve a high availability-like effect if duplicate polling is acceptable.

There are several common strategies used with hardware/software that provide high availability. Knowing these models and methods of redundancy (Tables 1 and 2) will help in mapping high availability priorities to available functions and capabilities.

Table 1. Models of Redundancy

Active/Active	Multiple application servers take an active role in polling, provisioning, or receiving alerts. There are no idle systems waiting to "take over," as each system is already doing practical work for a fractional part of the environment. In a dual-server Active/Active model each system should be configured with less than half of the total workload normally performed so that it can take the entire workload in case of a failure. All systems in this model maintain the state of the network and check for reachability with peer systems.
Active/Passive	Multiple application servers exist, but only one is taking an active role in polling, provisioning, or receiving alerts. The standby, or passive, system provides redundancy by coming online only when the primary system fails. The passive system receives periodic updates from the active server in order to stay current.

Table 2. Methods of Redundancy

Hot standby	Primary and secondary systems operate simultaneously. Data is replicated to the secondary server in real time. Both systems contain identical information.
Warm standby	The secondary system runs autonomously from the primary system. Data is replicated at scheduled periods. Minor data differences exist between scheduled replications.
Cold standby	The secondary system becomes operational only when the primary system fails. The cold standby receives scheduled replication updates.

High availability environments have unique requirements for managing licensing, application, and hardware capacity. Some advanced applications can handle more than dual-server redundancy. This option may be necessary for scalability, redundancy, or both. In this case, make sure the licensing, application, and server hardware capacity has adequate capacity for handling the additional load required in a failure mode.

Consider the following workload distributions scenarios:

- Two-server solution
  - First server managing to 70 percent of license, application, or server hardware capacity
  - Second server managing to 15 percent of capacity
  - A single server failure occurs
  - Second server takes 85 percent of the load in failure mode; acceptable
- Two-server solution
  - First server managing to 60 percent of capacity
  - Second server managing to 75 percent of capacity
  - A single server failure occurs

- Second server is not able to manage 135 percent of capacity

- Three-server solution (equal split)
  - First server managing to 40 percent of license, application, or server hardware capacity
  - Second server managing to 25 percent of license, application, or server hardware capacity
  - Third server managing to 35 percent of license, application, or server hardware capacity
  - A single server failure occurs—the 40 percent load needs to be split with each of the two remaining servers
  - Second server now manages 45 percent of capacity
  - Third server now manages 55 percent of capacity
  - This would be acceptable.
- Three-server solution (“all or nothing”)
  - First server managing to 40 percent of license, application, or server hardware capacity
  - Second server managing to 25 percent of license, application, or server hardware capacity
  - Third server managing to 35 percent of license, application, or server hardware capacity
  - A single server failure occurs—all 40 percent needs to go to one of the remaining servers
  - Second server now manages 65 percent of capacity
  - Third server continues to manage 35 percent of capacity
  - This would be acceptable.

In the most critical high availability scenarios an Active/Active with hot standby distribution is the preferred choice. For scenarios where high availability is needed, but financial considerations or product capabilities limit options, an Active/Passive with cold standby distribution may be warranted.

Cisco Network Management Product Examples

Cisco Applications and High Availability Functionality

Table 3 gives a short list of Cisco network management products and their high availability potential.

Table 3. Cisco Network Management Products and Their High Availability Potential

Application	Application/Database Reundancy	Caveats/Notes
CiscoWorks LAN Management Solution (LMS)	No	Device Credential Repository (DCR) can be spread across multiple servers to share managed device list and credentials – redundant server would be cold standby or duplicate polling
CiscoWorks Network Compliance Manager (NCM)	Optional	
CiscoWorks Unified Operations Manager	No	Similar as LMS (DCR) – cold standby or duplicate polling
CiscoWorks QoS Performance Manager (QPM)	No	Similar as QPM (DCR) – cold standby or duplicate polling
Cisco Info Center	Optional	
Active Network Abstraction (ANA)	Optional	Called “Protection Groups”

Cisco Network Registrar	Yes	Dynamic Host Configuration Protocol (DHCP) failover pairs, Domain Name System (DNS) secondary servers – local and regional clusters
CiscoSecure Access Controll Server	Yes	Database replication
Cisco IP Solution Center	No	No native high availability, but a customized solution is possible with Oracle, Veritas, and custom database synchronizing scripts
Cisco Application Networking Manager	Yes	

Optional – An additional feature license is necessary
Yes – feature is supported without additional cost

If a network management application does not support a true high availability model, it may still be possible to use a pair of autonomously running instances to achieve the same effect. Typically the concerns about running multiple network management applications for high availability are related to excessive network management polling. Combining the collected data into a single, cohesive instance or report is typically a concern when using multiple servers for scalability. In a high availability sense both servers run simultaneously, so both would have visibility to the same device information.

Using CiscoWorks LMS as an example, we can identify methods in using a non-high availability network management product in a high availability sense. Later in the document we will cover CiscoWorks NCM, which does have true high availability capabilities.

Running CiscoWorks LMS in a High Availability—Like Model

CiscoWorks LMS does not have true, native high availability functionality, with the notable exception being the Device Credential Repository (DCR) component of Common Services. DCR synchronizes the managed device list and the credentials across the applications within the CiscoWorks LMS suite. Individual applications, such as Resource Manager Essentials, Campus Manager, and so on, do not have data sharing capability across multiple instances.

To run CiscoWorks LMS in a high availability-like model requires two servers and two licensed copies of CiscoWorks LMS. One server is considered the primary server and the second server the secondary server. The secondary server is configured to poll the managed network devices at a rate four to five times slower than the primary, which is running at default polling intervals. The failover from primary to secondary is a manual process. It is important to have application and server monitoring solutions in place to notify the user if the primary server is down so the manual process for failover can be initiated.

The LMS application administrator should be familiar with how to adjust the various application polling settings quickly in case failover is required. One method would be through the CiscoWorks LMS 3.0 Setup Center and the Server Settings option (Figure 4).

Figure 4. CiscoWorks LMS Setup Center

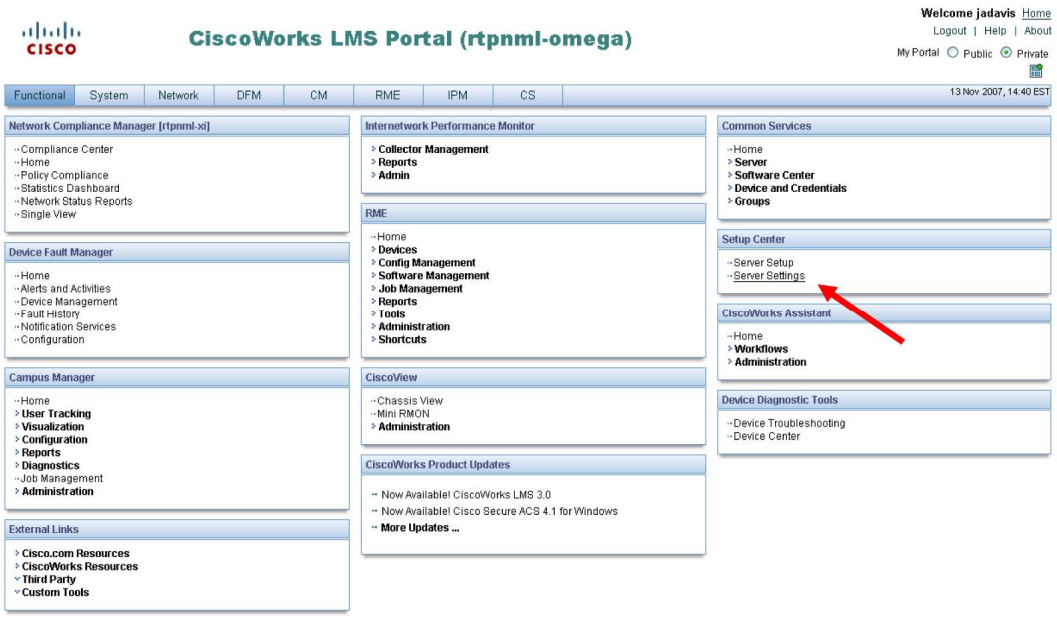

Modifying the **Data Collection Settings** and **Data Collection Schedule** menu options would be the quickest way to adjust from four to five times slower polling back to normal rates while the primary server is unavailable (Figures 5 and 6).

Figure 5. CiscoWorks LMS Setup Center

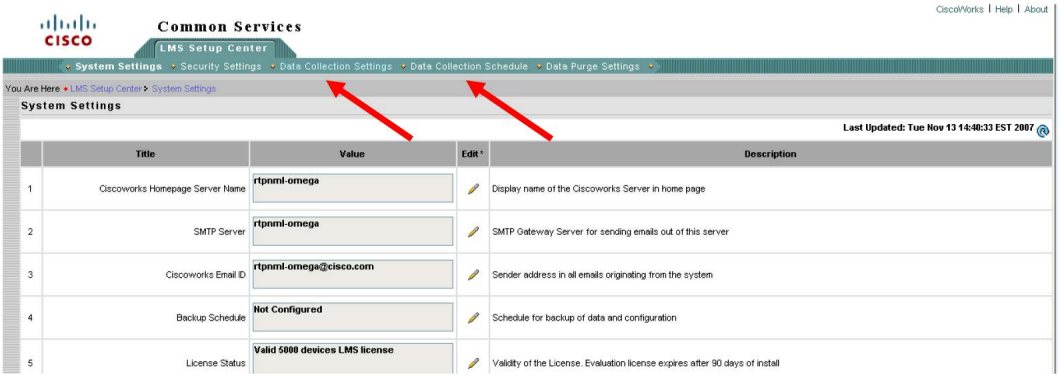

Figure 6. CiscoWorks LMS Setup Center (Continued)

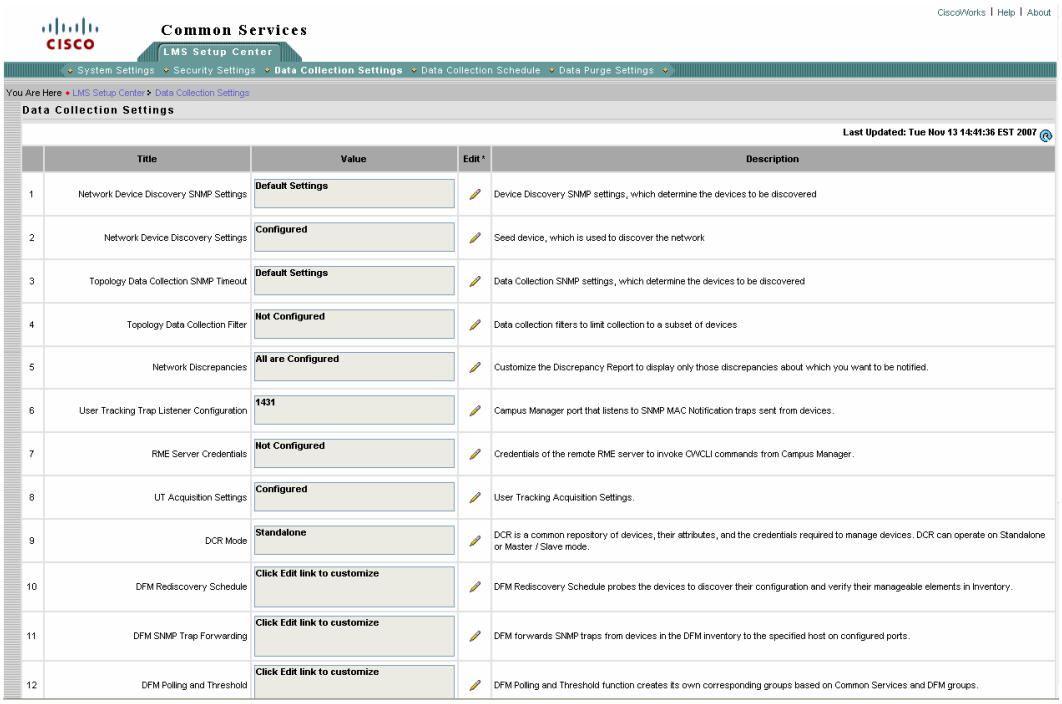

The managed network devices would be configured to point their syslog event messages and SNMP traps, if Device Fault Manager is being used, to both servers simultaneously. In this manner, both servers maintain near real-time updates to configuration and hardware/inventory changes. Both servers receive syslog event messages and would be triggered to poll the device for the latest change.

It has been observed in several environments that even three or four CiscoWorks LMS servers would not place undue burden on the managed devices. If the device configurations are especially long and the managed device does not respond appropriately to multiple requests for SNMP-initiated TFTP sessions for the configuration file, set the secondary server to use a different transport protocol, other than SNMP, to obtain the configuration file (Figures 7 and 8).

Figure 7. Updating CiscoWorks Resource Manager Essentials Configuration Management Transport Protocols

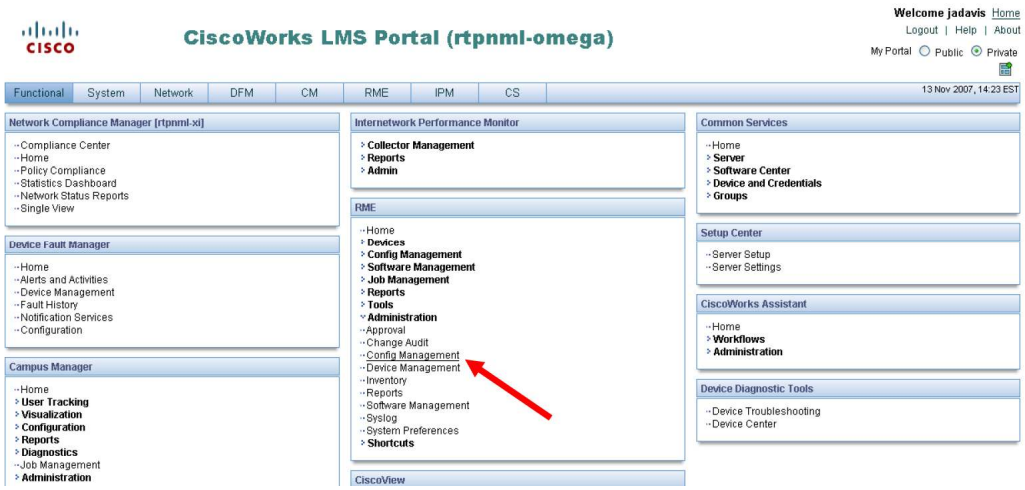

Figure 8. Updating CiscoWorks Resource Manager Essentials Configuration Management Transport Protocols (Continued)

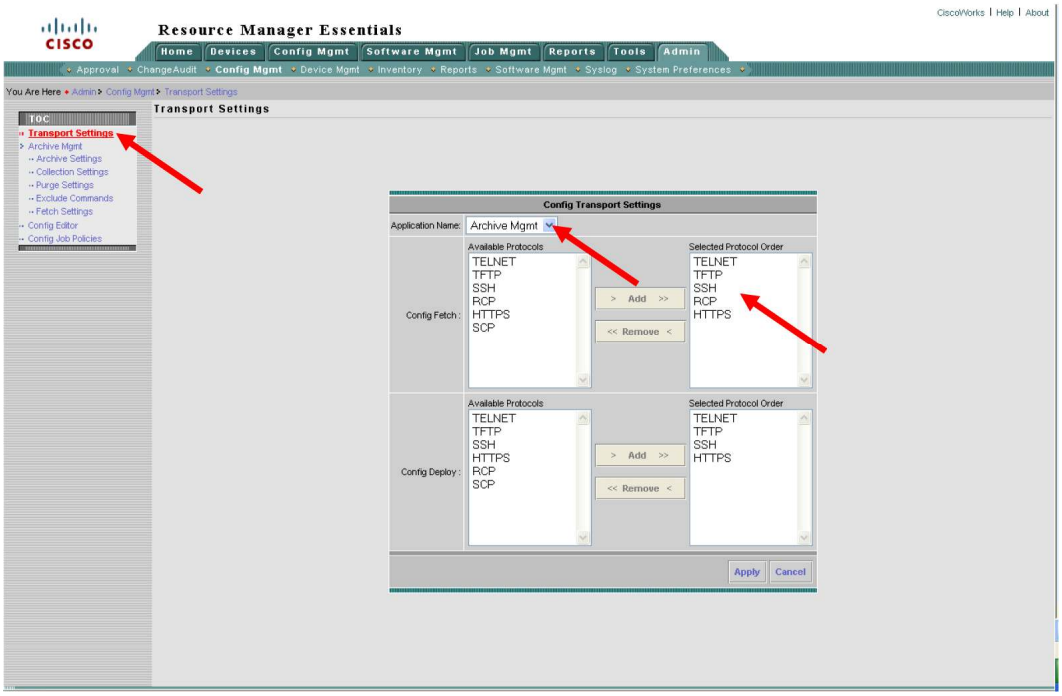

With both servers configured for the same device list, both servers receiving syslog event messages (and SNMP traps, if DFM is used) and the secondary server set to poll at a four to five times slower rate, then both systems will maintain a reasonable state of device configuration, inventory, and fault status. If the primary system has a failure, the secondary is ready with simple polling tune-ups.

Running CiscoWorks NCM for High Availability

Another popular tool that Cisco offers is CiscoWorks Network Compliance Manager (NCM). It does have native high availability capabilities that can be licensed. Refer to the configuration guide at http://www.cisco.com/en/US/docs/net_mgmt/ciscoworks_network_compliance_manager/1.1/high_availability/configuration/guide/HA_guide.pdf

Implementing NCM high availability requires the use of an Oracle or Microsoft SQL Server database.

It is fine to use Oracle or Microsoft SQL Server on the same server as the NCM application up until 2,000 devices are managed. Past that point, there will need to be separate application and database servers. In a high availability framework, it would obviously mean two application servers and two database servers.

NCM supports greater than dual-server redundancy for scalability or fault tolerance, so the previous comments about managing within appropriate licensing, application, and hardware capacities should be noted.

The NCM high availability framework configures the database servers to continually synchronize configuration, polling, and change-related information. Each NCM application server can be the primary poller for devices closest to it and secondary for devices in another area. In this way the network management work-load is spread across the servers (Figure 9).

Figure 9. CiscoWorks NCM Database Synchronization Framework

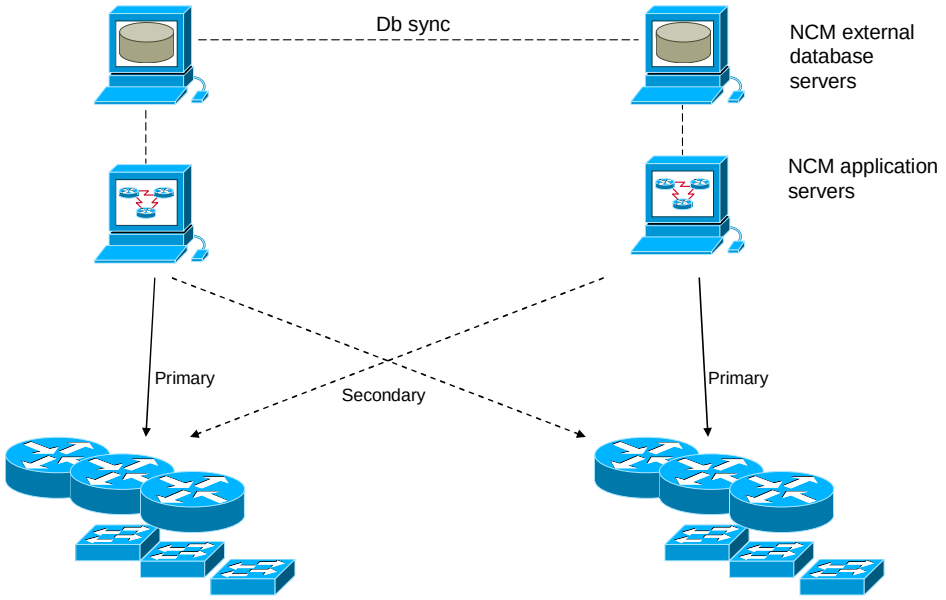

The application server checks the health of its paired database (or external database server). When a failure occurs and the high availability feature is needed, a notification is sent by e-mail; the NCM administrator would log into the surviving NCM application server and move the polling responsibilities for the failed server domain to it.

Running CiscoSecure ACS for High Availability

CiscoSecure Access Control Server (CS-ACS) can be run in a highly available scenario with database replication among the ACS servers. One server is selected as the primary. Ideally no AAA protocol operations are performed from the managed devices to this server. ACS administration should be the only tasks performed on this system. A number of subordinate ACS servers are configured to be replication peers with the primary or with each other. The number of subordinate ACS servers is dependent upon AAA transaction load and scale of the network being served. For assistance in determining how many ACS servers (or appliances) to deploy, refer to http://www.cisco.com/en/US/docs/net_mgmt/cisco_secure_access_control_server_for_windows/4.1/configuration/guide/deploy.html#wp1041903

All ACS servers must be at the same software version for replication. Configuring ACS replication is covered in the following online documents:
http://www.cisco.com/en/US/docs/net_mgmt/cisco_secure_access_control_server_for_windows/4.1/configuration/guide/deploy.html#wp1049439
http://www.cisco.com/en/US/docs/net_mgmt/cisco_secure_access_control_server_for_windows/4.1/user/SCAdv.html#wp755988

An important design consideration to take into account is that the ACS authentication services are suspended briefly when a server is preparing or importing the database. The managed network devices should have multiple ACS servers defined, and the ACS server topology should be arranged in a cascaded fashion. In this manner there will always be an ACS server that the managed device can reach for AAA services (Figure 10).

Figure 10. Example 1 – Poor Three-Server Design

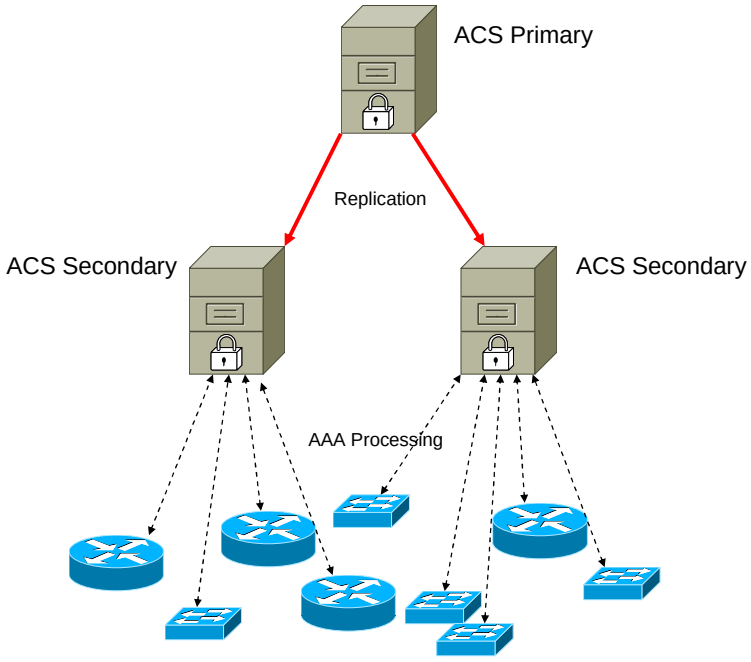

In this scenario replication was configured with the secondary servers in parallel. Both secondaries would have their authentication processes suspended simultaneously for a brief period when replication happens. Since the managed devices only have one ACS server configured, they would have no AAA service during the brief replication process.

In this scenario, even if the managed devices had both ACS secondaries configured, the devices would not have AAA services during replication. The next design scenario is best for a three-server environment (Figure 11).

Figure 11. Example 2 – Good Three-Server Design

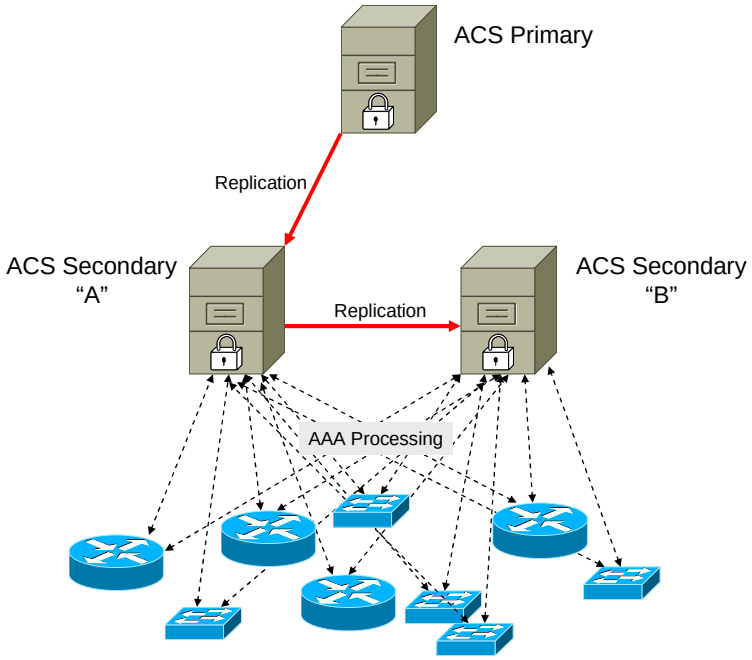

In this scenario replication happens from the primary to secondary “A”, then from secondary “A” to secondary “B”. At all times there is at least one ACS server operational. Since all managed devices have dual ACS server configurations, they will continue AAA processing when “A” or “B” briefly suspend authentication for replication.

Larger ACS frameworks are certainly possible. It is a leading practice to use a cascaded design for replication. Parallel cascaded designs would offer even more resiliency for AAA services. Adding a third ACS server to the managed device configurations would also enhance AAA service resiliency.

CiscoWorks Active Network Abstraction (ANA) and High Availability

Cisco ANA is a new platform that facilitates scalable, vendor-neutral network resource management in a multitechnology, multiservice environment. It provides network element configuration and fault management. It uses a near real-time virtual network model to provide an open, standards-based approach to manipulating devices through APIs that abstract the real network. ANA does provide a rich high availability capability. ANA is a high-complexity and broadly customizable platform that requires implementation assistance from Cisco Advanced Services.

The ANA unit servers represent the network elements virtually and provide communication to the physical network elements and ANA gateways. The application provides high availability in the unit servers through a feature that allows the unit servers to be clustered in protection groups. Protection groups operate so that if a unit server becomes unavailable another unit server will take over the load.

The ANA gateways provide the user and North Bound Interface (NBI) application connectivity to the network. High availability for the gateway servers is provided through the design and implementation of the Oracle database, Veritas, and hardware RAID on the servers (Figure 12).

Figure 12. Cisco ANA Architecture for High Availability

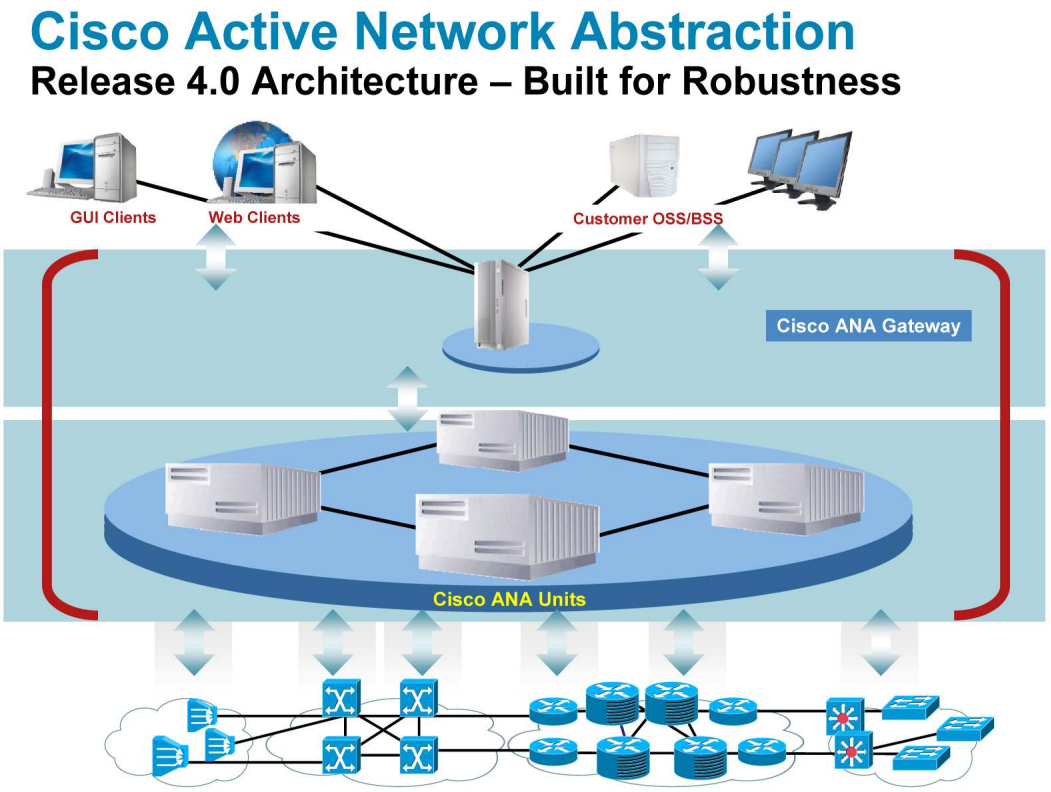

For more background on ANA, refer to http://www.cisco.com/go/ana

For more information on ANA installation, refer to http://www.cisco.com/en/US/docs/net_mgmt/active_network_abstraction/4.0/installation/guide/install.html

Cisco IP Solution Center and High Availability

Cisco IP Solution Center (ISC) is a network management solution for resource management and profile-based provisioning of Multiprotocol Label Switching (MPLS), Metro Ethernet networks and tunnel engineering.

ISC does not provide a native high availability capability. High availability is achieved through careful design and implementation of redundant Oracle databases and hardware RAID with database synchronization scripts.

The design consultation and installation services are offered through Cisco Advanced Services. This form of high availability would be considered Active/Passive with warm standby (Figure 13).

Figure 13. ISC Architecture for High Availability

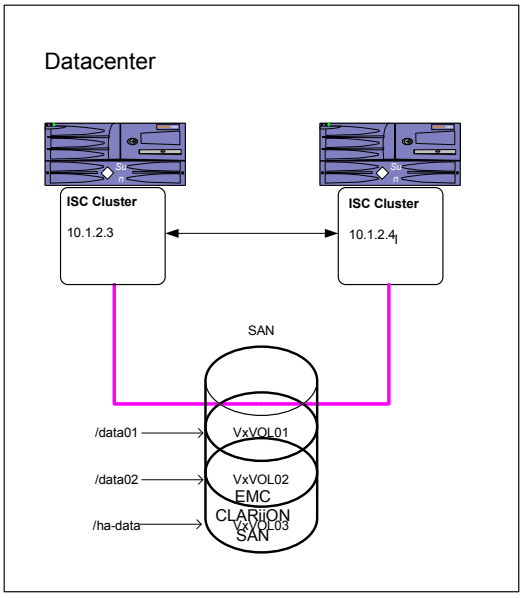

For more information on ISC installation, refer to http://www.cisco.com/en/US/products/sw/netmgtsw/ps4748/products_installation_guide_book09186a008081d37b.html

Conclusions

Running network management applications in a highly available fashion is required in many environments. Identifying the key network management functions and the application capabilities will drive your high availability considerations. The methods described earlier can be used to make non-high availability applications operate in a high availability-like model.

For additional assistance with network management services, please contact your Cisco service account manager for engagement with Cisco Advanced Services.

Acronyms

Acronyms	Definition
NMS	Network management system
NCM	Network Compliance Manager – a management product in the CiscoWorks family
DCR	Device Credential Repository – a CiscoWorks Common Services component that synchronizes device lists and credentials across multiple CiscoWorks servers
AAA	Authentication, authorization, and accounting
NBI	Northbound Interface – provides an output-only interface to higher levels in an architecture (for example, syslog)
ACS	Access Control Server – a CiscoSecure product family offering that performs network AAA services

References

[1] Balabit Syslog-NG: http://www.balabit.com/network-security/syslog-ng/

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
168 Robinson Road
#28-01 Capital Tower
Singapore 068912
www.cisco.com
Tel: +65 6317 7777
Fax: +65 6317 7799

Europe Headquarters
Cisco Systems International BV
Haarlerbergpark
Haarlerbergweg 13-19
1101 CH Amsterdam
The Netherlands
www-europe.cisco.com
Tel: +31 0 800 020 0791
Fax: +31 0 20 357 1100

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

Printed in USA C11-449655-00 1/08