Network and Scalability Whitepaper

Resource Management Suite - RMS Enterprise Software

Centralized remote management of networked AV equipment and building systems

The software features a user-friendly dashboard making it easy to centralize the management and monitoring of AV equipment, lights, HVAC and other building functions. RMS Enterprise contributes to energy reduction initiatives and extending useful life of devices. Additionally, IT and AV managers can improve SLA response times and reduce equipment downtime.

It's Your World. Take Control.

Table of Contents

It's Your World. Take Control.

Introduction

AMX's Resource Management Suite (RMS) is designed for conference rooms, auditoriums and data centers where IT and AV managers can benefit from a centralized remote management tool with multiple capabilities to monitor and manage equipment, automate meeting room setup based on a user ID, or provide reports on energy usage and maintenance schedules.

Information Technology managers are now making AV decisions with enterprise requirements such as…

- Traditional client-server model using standard ports and communications
- Virtual servers, hosted environments, redundancy, and server farm deployments
- Large scale management for users, roles, and, tracking user activities

Texas Woman's University Network Operation Center – Denton, Tx

RMS Enterprise meets those requirements and offers many features that benefit both AV and IT managers.

It's Your World. Take Control.

Enterprise Architecture

RMS Enterprise provides large scale management for user and roles, as well for tracking user activities including an audit trail of who performed each activity and when it was completed. The server software supports authentication, encryption and protection from cross-site scripting to prevent security threats**.** The use of Hibernate and its parameterized queries protects RMS server against SQL injection attacks.

RMS Enterprise supports globally deployed systems by adding support for WAN and Internet based connections. A client-server communication strategy utilizing a service-oriented architecture (SOA), supports globally distributed client/endpoint communication. RMS Enterprise is a client/server application where the NetLinx system acts as the client and the RMS Enterprise application server listens for connections from NetLinx systems. NetLinx and the RMS Enterprise application server communicate using TCP/IP sockets.

In order to establish communication, each NetLinx system must be able to resolve and connect to the RMS Enterprise application server. This can be accomplished with a variety of Network configurations including local area networks (LAN), wide area networks (WAN), and the Internet.

LDAP User Management

RMS Enterprise supports Directory service based user management allowing authentication, authorization, group/permission association, and synchronization with an LDAP or Windows Active Directory server. The Lightweight Directory Access Protocol (LDAP) is used to read from and write to Active Directory. By default, LDAP traffic is transmitted unsecured. You can make LDAP traffic confidential and secure by using Secure Sockets Layer (SSL) / Transport Layer Security (TLS) technology. You can enable LDAP over SSL (LDAPS) by installing a properly formatted certificate from either a Microsoft certification authority (CA) or a non-Microsoft CA.

AMX
3000 Research Drive
Richardson, TX 75082

tel 800.222.0193
fax 469.624.7153

www.amx.com

It's Your World. Take Control.

4 | P a g e

Audit Trail for Activities

RMS Enterprise provides an audit trail for user-initiated activities. This audit tracking system will aid in tracing a user's interaction with the RMS system. The auditable items for RMS are limited to user invoked actions including timestamp when the action took place, a description of the action that took place, name of the user invoking the action, subject of the action, and location where an action took place.

Service Oriented Architecture

RMS Enterprise is a traditional client-server model where endpoint RMS SDK enabled systems act as clients communicating with the central RMS services. This implementation provides well defined interfaces that are easily consumable by other client platforms/systems. The SOA implementation utilizes web friendly communication ports and conventions that are more readily adopted and understood by IT organizations.

It's Your World. Take Control.

Enterprise Scaled Deployments
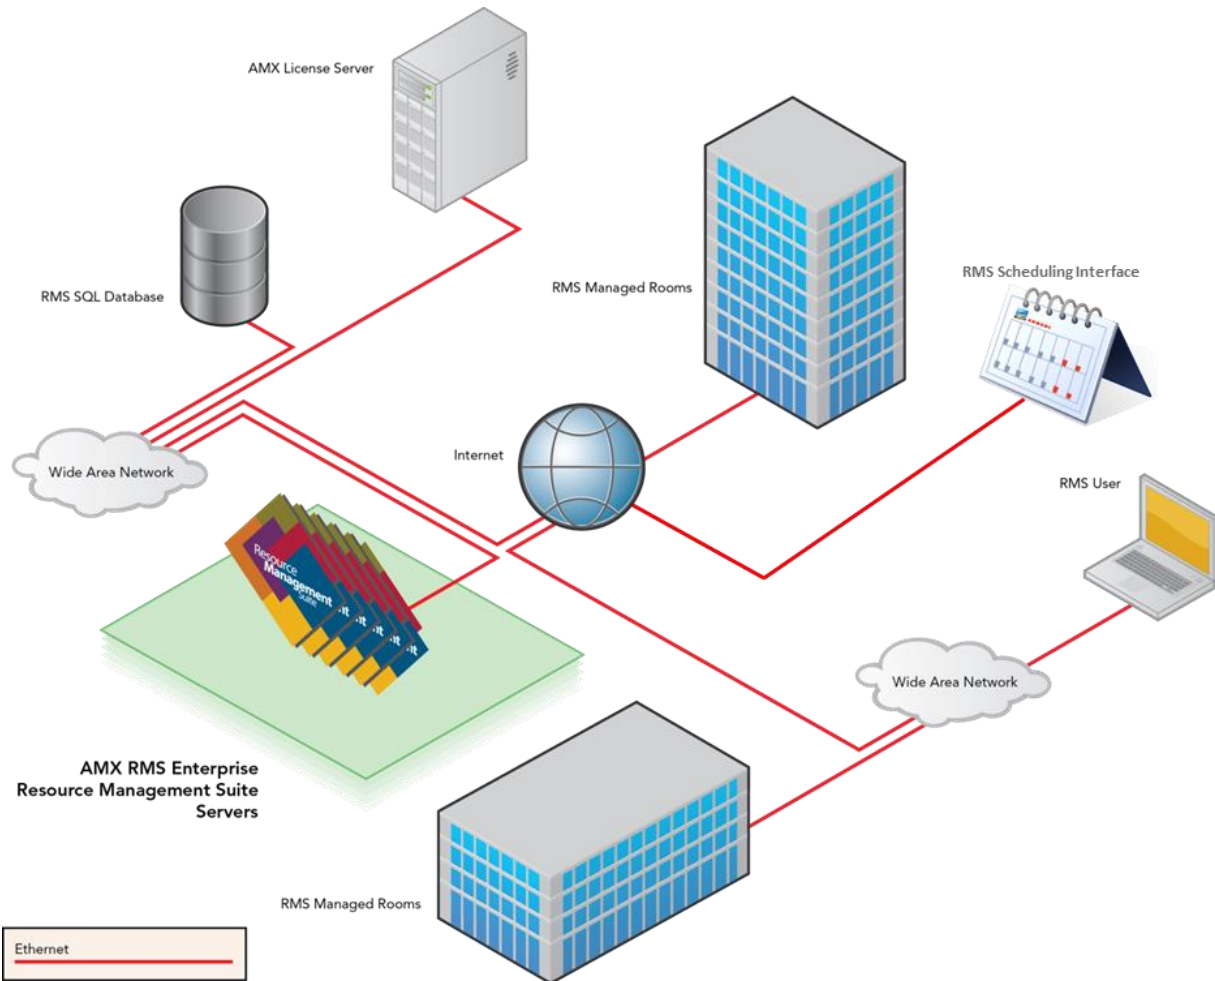
RMS Enterprise may be deployed in a number of configurations to satisfy the needs of the customer. These deployment options include a single server stand-alone solution in a local network up to a multi-server deployment in a web farm for scalability, redundancy, and load balancing.
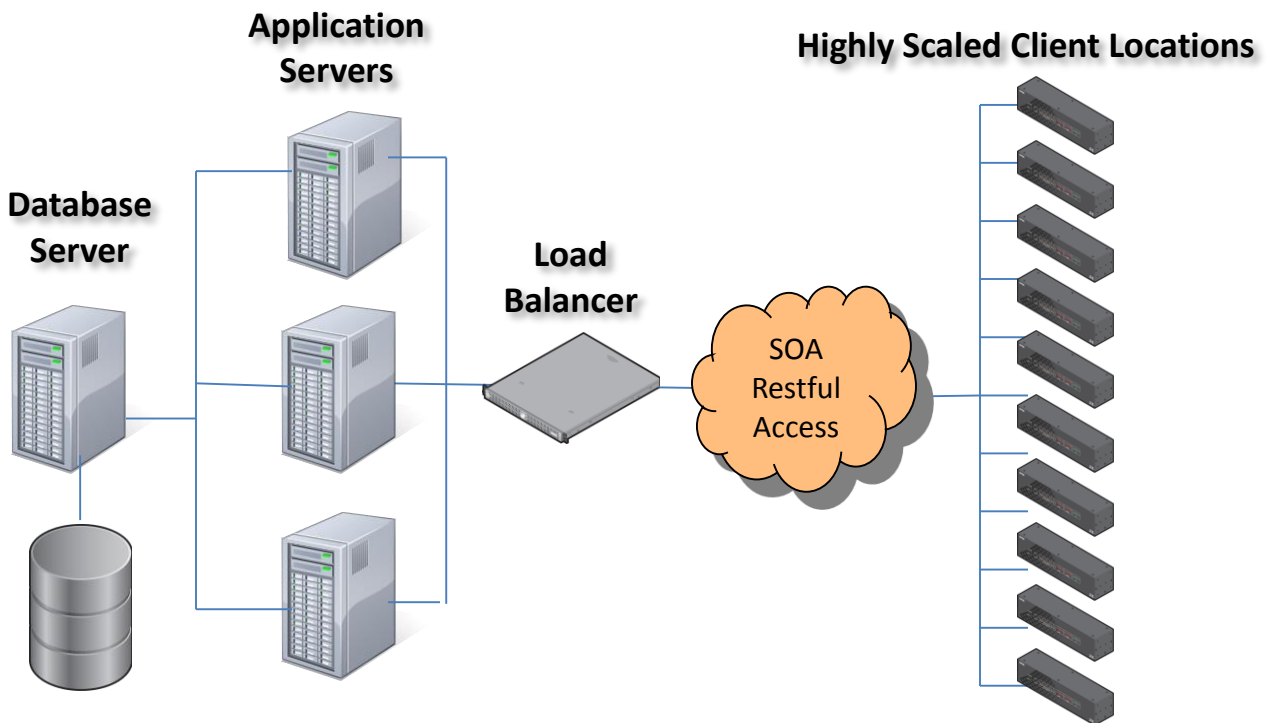
It's Your World. Take Control.

RMS Scaling Architecture

RMS Enterprise is deployed as a client - server based application communicating across private and/or public networks. RMS connectivity and communication requires a TCP/IP network with a resolvable server address by all client endpoints. The RMS server handles all communication with client endpoints via common protocols (HTTP) and warehouses all data in an SQL database ensuring it can always cleanly and efficiently integrate with the rest of your infrastructure and scale with your needs.

Application Servers

Highly Scaled Client Locations

Database Server

Load Balancer

SOA Restful Access

It's Your World. Take Control.

RMS Enterprise Multi-Server Deployments

RMS Enterprise provides a scalable application service provided by multi-server / server cluster deployment. The RMS services are designed to work across multiple servers in a server cluster infrastructure and on a single server deployment. Taking advantage of a web farm infrastructure can provide greater scalability, service redundancy and traffic load balancing. Special consideration must be taken when determining data and user file storage such that the data stores are accessible from all servers. The RMS database should be used as the common data store. The RMS Client SDK needs no specific knowledge or configuration to take advantage of the multi-server infrastructure. RMS Enterprise performs best when installed on a server that is exclusively dedicated to running the RMS Enterprise application. For optimal performance, it is strongly recommended that a dedicated RMS server is used to run the RMS Enterprise application - regardless of the size of the installation. RMS Enterprise will easily handle larger deployments however to guarantee system performance it is important to ensure that when deploying to environments with more than 50 systems the host machine (whether physical or virtual) be dedicated to this application.

Supported Virtual Servers

RMS Enterprise can also be hosted on a virtual server with either of the following virtualization systems:

- VMware ESXi (4.x or greater)
- Hyper-V

If RMS Enterprise is to be installed on a virtual server, verify that it has sufficient hardware resources to support these virtual hosts. The created virtual servers must meet the Minimum Server Hardware Requirements indicated above.

Server Software Requirements

Please reference the RMS Enterprise Installation Guide for the latest server and software requirements.

http://www.amx.com//assets/manuals/RMSEnterprise.InstallationGuide.pdf
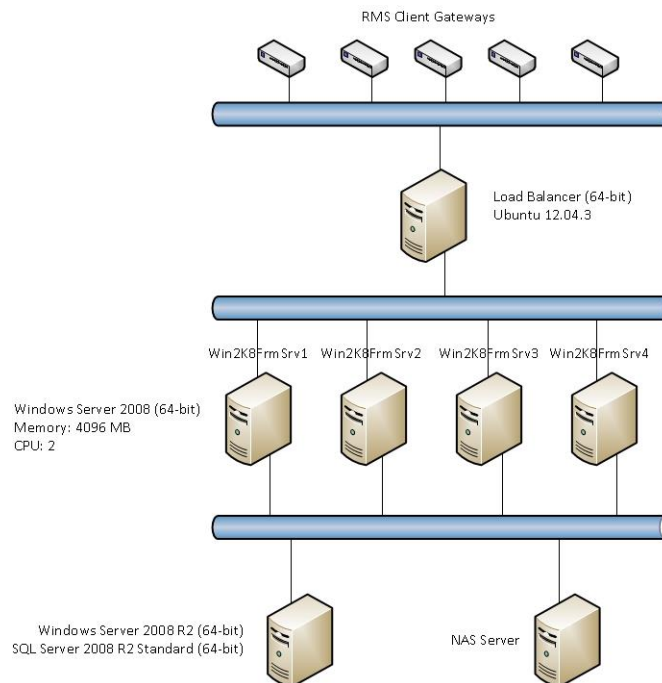
It's Your World. Take Control.

Multi Server – Best Practices

This section is designed to guide the user for best practices for setting up Resource Management Suite (RMS) servers in a multi-server or virtual environment. Note that Multi-Server Support is disabled by default. The options in this dialog are only available if the Enable Multi-Server option is selected.

Apache web server was used to perform load balancing to the backend RMS servers which run in Tomcat. Apache web server can be installed on Linux, BSD or Windows. Screen captures in this document are taken from the Ubuntu 12 Linux distro.

AMX tested 8000 simulated client gateways across four RMS servers deployed in a multi-server configuration. Each of the servers ideally had 2000 client gateways, when a server went down, its traffic would be redirected and split across the remaining three servers. Each of the RMS servers are identical in hardware configuration.

VMware ESXi

Following are some additional configurations for each of the virtual machines that were setup in the AMX tested multi-server environment. For each of the virtual machines, right click on the desired VM and select `Edit Settings…` from the context menu to edit the settings, then select the 'Resources' tab.
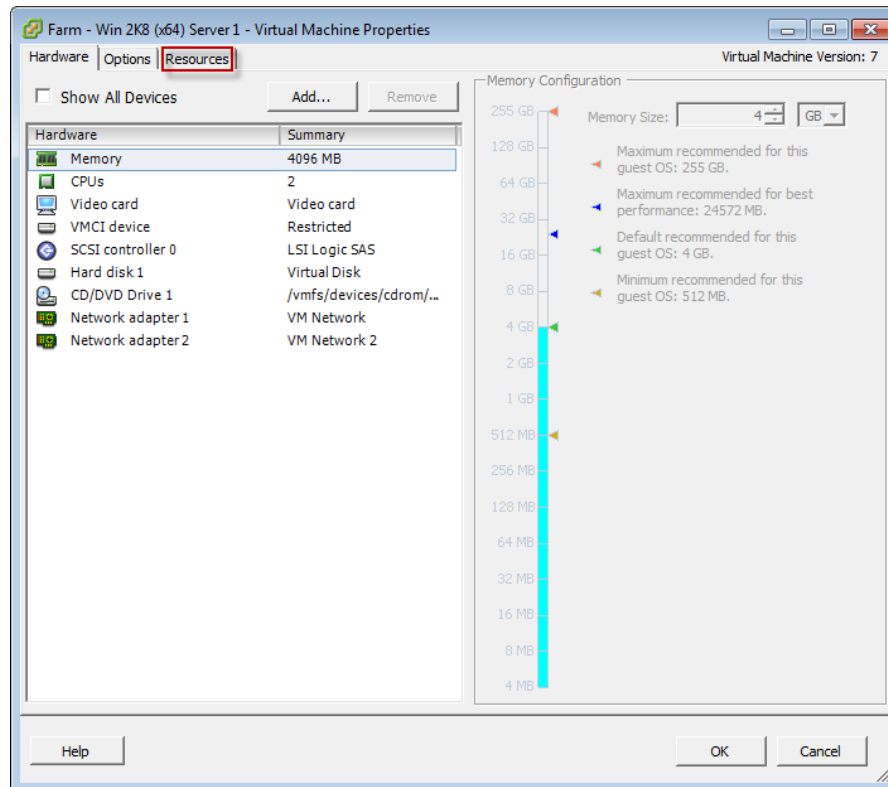
First configure the CPU resources, this lets you allocate processor resources for a virtual machine, specifying reservations, limits, and shares. You must power off the virtual machine before configuring CPU resources.

It's Your World. Take Control.

Under `Resource Allocation`, the `Shares` configures the CPU shares for this virtual machine with respect to the parent's total. Sibling virtual machines share resources according to their relative share values bounded by the reservation and limit – select `High`.

Under `Resource Allocation`, the `Reservation` configures the upper limit for this virtual machine's CPU allocation. Move the slider as far as it will go to the right.
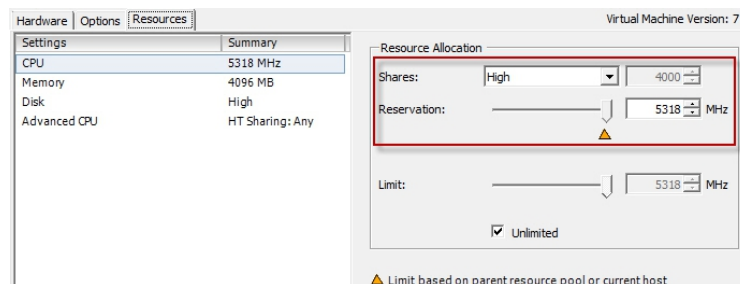
The memory resources page lets you allocate memory resources for a virtual machine, specifying reservations, limits, and shares. Memory resource configuration is distinct from the virtual hardware memory size, which is configured on the `Hardware` tab. The memory resource setting determines how much of the host's memory is allocated to a virtual machine. The virtual hardware memory size determines how much memory applications that are running inside the virtual machine have available to them. A virtual machine cannot benefit from more memory resources than its configured virtual hardware memory size. You must power off the virtual machine before configuring memory resources.

Under `Resource Allocation`, the `Shares` configures memory shares for this virtual machine with respect to the parent's total. Sibling virtual machines share resources according to their relative share values bounded by the reservation and limit – select `High`.
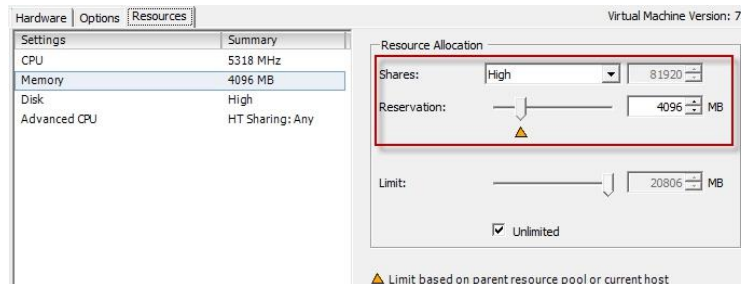
Under `Resource Allocation`, the `Reservation` configures the guaranteed memory allocation for this virtual machine, slide this all the way to the right.

It's Your World. Take Control.

The disk resources page lets you allocate host disk I/O bandwidth to the virtual hard disks of a virtual machine. Note: In general, CPU and memory resources are more likely to constrain virtual machine performance than disk resources.

`Shares` represent the relative importance of a virtual machine with regard to the distribution of storage I/O resources. Virtual machines with higher share values allow more through-put and have lower latency. Disk bandwidth shares for this disk on this virtual machine with respect to all disks on all virtual machines for each datastore that the virtual machine accesses. Disks share resources according to their relative share values.

Click the `Shares` field and change the value to allocate a number of shares of its disk bandwidth to the virtual machine – select `High` (2000). Verify that the `Limit` is set to the default value of `unlimited`.

SQL Server Configuration

The SQL server should have several available drives, in the case of the AMX test environment, the server had three hard drives on different controllers:

1. [C:\] Windows Server 2008 OS and SQL Server 2008 R2 Standard Edition

It's Your World. Take Control.

2. [E:\] Database default location for log files
3. [F:\] Database default location for data Files

This also needs to be configured in the SQL Server, to do that, start the *Microsoft SQL Server Management Studio* application and select the main node in the *Object Explorer* frame and select `Properties` from the context menu.
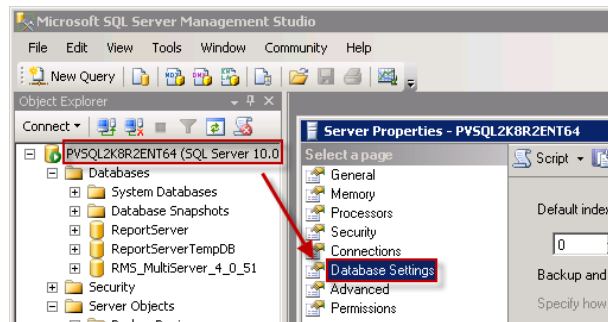
Under the server properties, select the `Database Setting` option and update the database default locations to point to the new drives. This should help in the SQL performance.

While editing the properties for the SQL server, one additional adjustment should be to the SQL server memory. SQL server can (and will) starve the operating system of memory, so it is advisable to set the maximum server memory that it can use. Figure out how much memory you need for the OS and set the rest of the memory for the SQL Server; in the AMX test case, the server had 22GB so we assigned the maximum memory the SQL server could use to 20GB.

AMX
3000 Research Drive
Richardson, TX 75082

tel 800.222.0193
fax 469.624.7153

www.amx.com

It's Your World. Take Control.

13 | P a g e
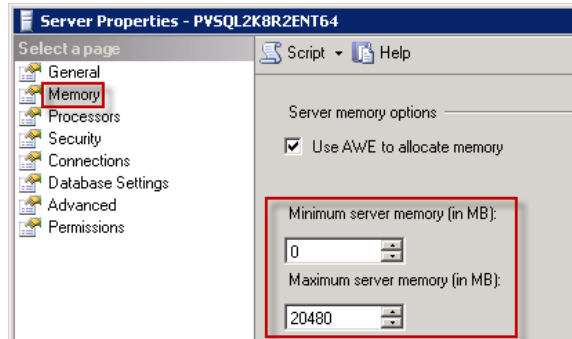
Apache Configuration (Load Balancer)

Apache will need to serve many concurrent requests and therefore must be configured to handle a lot of traffic. Use the Multi-Processing modules for scaling Apache.

RMS requires sticky sessions for the UI. Sticky sessions for the client is preferred, but not strictly required. For simplicity, it is much easier to simply use sticky sessions for both.

Finally, the AJP protocol is being used for communication between Apache and Tomcat. It is possible for Apache to simply forward requests onto Tomcat over HTTP, but that was not the configuration we opted for.

Multi-Processing Module

Use the mpm-worker[1] module for scalability in Apache itself. This has to be built into the Apache binary. By default, Apache usually comes with "prefork". To install the version of Apache with mpm-worker compiled in, do this via the command prompt (Ubuntu only). To access the command prompt you can ssh into the load balancer or just open a console via the vSphere Client:

```
sudo apt-get install apache2-mpm-worker
```

---

[1] http://httpd.apache.org/docs/2.2/mod/worker.html

It's Your World. Take Control.

Here is AMX's snippet of our `/etc/apache2/apache.conf` file for the configuration of mpm-worker:

```
<IfModule mpm_worker_module>
    StartServers         20
    ServerLimit          240
    MinSpareThreads      25
    MaxSpareThreads      75
    ThreadLimit          75
    ThreadsPerChild      50
    MaxClients           12000
    ListenBacklog        2000
    MaxRequestsPerChild  0
</IfModule>
```

Apache Modules

The following modules need to be enabled if they aren't by default. Use a2enmod to enable them in Ubuntu.

- headers   (*needed for sticky session support via request headers*)
- proxy
- proxy_ajp   (*AJP is the protocol used between Apache and Tomcat*)
- proxy_balancer   (*module for load-balancing*)
- proxy_connect

```
sudo a2enmod headers proxy proxy_ajp proxy_balancer proxy_connect
```

[1] http://httpd.apache.org/docs/2.2/mod/worker.html

Load Balancer

PV added the keepalive=on parameter in hopes that this will assist with the load balancer error.[2] Here is what PV's `/etc/apache2/httpd.conf` file looks like:

```
<Proxy balancer://rmscluster>
  BalancerMember ajp://Win2k8FrmSrv1:8009 route=Win2k8FrmSrv1 loadfactor=1 max=3000 keepalive=on
  BalancerMember ajp://Win2k8FrmSrv2:8009 route=Win2k8FrmSrv2 loadfactor=1 max=3000 keepalive=on
  BalancerMember ajp://Win2k8FrmSrv3:8009 route=Win2k8FrmSrv3 loadfactor=1 max=3000 keepalive=on
  BalancerMember ajp://Win2k8FrmSrv4:8009 route=Win2k8FrmSrv4 loadfactor=1 max=3000 keepalive=on

  ProxySet lbmethod=byrequests stickysession=JSESSIONID|jsessionid scolonpathdelim=On
nofailover=Off
</Proxy>

ProxyRequests Off
ProxyPass /rms balancer://rmscluster/rms
ProxyPassReverse /rms balancer://rmscluster/rms

<Location /balancer-manager>
  SetHandler balancer-manager
  Order deny,allow
  Allow from all
</Location>
```

Please note that AMX software development found the $max=XYZ$ parameter on the `BalancerMember` definition to be very important in not sending too many concurrent requests to that backend. Also, the sum of all the max parameters should be at least as big as the `MaxClients` configured in the mpm-worker module.

Increasing ulimit

The *ulimit* is a mechanism for restricting the amount of various resources a process can consume. The following changes were made to increase the number of file descriptors for Apache from 8192.

/etc/apache2/envars

---

[2] (104)Connection reset by peer: ajp_ilink_receive() can't receive header
ajp_read_header ajp_ilink_receive
(120006)APR does not understand this error code: proxy: read response failed from *server*

It's Your World. Take Control.

The following change was made to the end of the file:

```
## If you need a higher file descriptor limit, uncomment and adjust the
## following line (default is 8192):
#APACHE_ULIMIT_MAX_FILES='ulimit -n 65536'
APACHE_ULIMIT_MAX_FILES='ulimit -n 16384'
```

/usr/sbin/apachectl

The following change was made, it is located about half-way down the file:

```
# Set this variable to a command that increases the maximum
# number of file descriptors allowed per child process. This is
# critical for configurations that use many file descriptors,
# such as mass vhosting, or a multithreaded server.
ULIMIT_MAX_FILES="${APACHE_ULIMIT_MAX_FILES:-ulimit -n 16384}"
```

[1] (104)Connection reset by peer: ajp_ilink_receive() can't receive header
ajp_read_header ajp_ilink_receive
        (120006)APR does not understand this error code: proxy: read response failed from *server*

Tomcat Configuration (RMS Servers)

All the following steps assume that the RMS Server has already been installed. The following Tomcat configuration changes need to be made to the `server.xml` file located in the `TOMCAT_HOME/conf` directory.

Thread Pool

For the best performance it is a good idea to configure a thread pool in Tomcat. This allows Tomcat to manage the number of threads, growing or shrinking the pool size based upon the work load. The optimum number of `maxThreads` for the pool is difficult to determine and may require some tuning. Certainly a one-to-one ratio of RMS client gateways to Tomcat threads is not necessary. Whatever number is chosen, it should not be smaller than the max value that was configured for the `BalancerMember` for this server in the Apache configuration (see above).

For our tests, the following thread pool was used on the larger server "perftest1" which served approximately 3000 client gateways:

```
<Executor name="tomcatThreadPool"
```
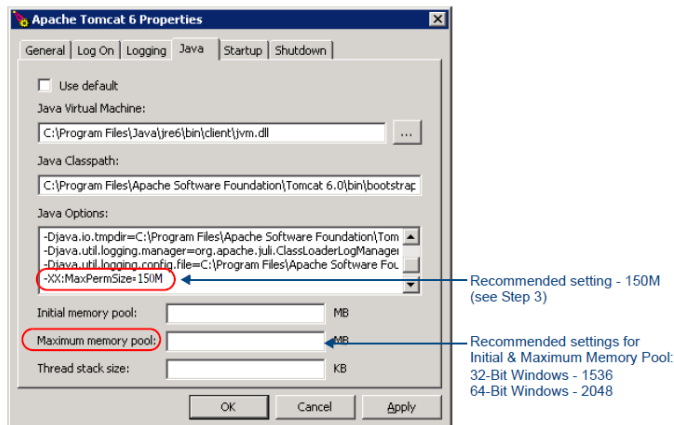
It's Your World. Take Control.

```
namePrefix="catalina-exec-"
maxThreads="1500"
minSpareThreads="50"/>
```

The following recommendations are for configuration of an average RMS server: In the Java tab, click inside the Java Options text box and scroll down to the last line. After the last line of code, enter the following line of code:

-XX:MaxPermSize=150M

In most cases, the Initial Memory Pool and Maximum Memory Pool values should be equal. The ideal values for these settings will depend on the number of clients connected to your RMS server. As a rule of thumb, consider allocating half of the server's available memory to Tomcat for use with RMS.The following recommendations apply for the Maximum Memory Pool value, according to the version of Windows (32-bit or 64-bit) that is installed on your server PC:

- 32-bit Windows - 1500MB (1.5GB)
- 64-bit Windows - 2048MB (2GB)

AJP Support

We used the AJP protocol between the Apache web server and the Tomcat application servers. Ensure the Connector for the AJP protocol is uncommented, and that it references the thread pool configured above.

It's Your World. Take Control.

```
<Connector executor="tomcatThreadPool"
          port="8009"
          protocol="AJP/1.3"
          connectionTimeout="60000"
          redirectPort="8443"/>
```

In addition, for load-balancing over AJP to function, the `jvmRoute` property needs to be set up. This name must match route defined on the `BalancerMember` in the Apache configuration (see above).

```
<Engine name="Win2k8FrmSrv1"
        defaultHost="localhost"
        jvmRoute="Win2k8FrmSrv1"/>
```

Windows Service Account for Tomcat

By default, Tomcat as a service runs under a "Local System account." Development had trouble getting this account to work properly with the Windows shared folder. The solution was to create local Windows accounts on each box that use the same user name ("rmsadmin") and password. Tomcat needs to be configured to use this account when running as a service.

The Tomcat Configuration GUI tool makes it simple to switch from Local System account to your own account.

It's Your World. Take Control.
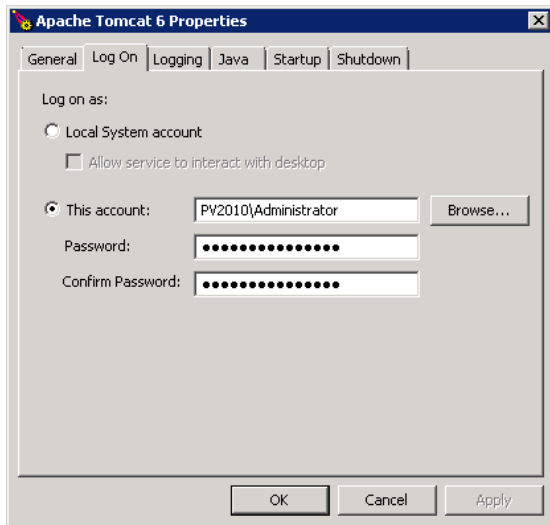
One side effect of creating a special Windows account is that by default Windows may not allow that account to run programs as services. Use the Windows tool `secpol.msc` to give your account the necessary permission. Once this tool is launched navigate to Local Policies > User Rights Assignment > Log on as a service, and then add your new user account to that security policy.
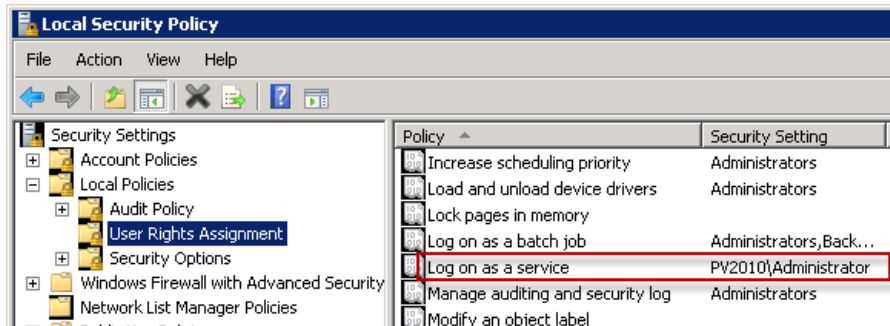
It's Your World. Take Control.

RMS Multi-Server Configuration

Licensing

The AMX License Manager is used to install and manage software licenses for RMS Enterprise (as well as other AMX software applications). The AMX License Manager handles two distinct aspects of the RMS Enterprise installation:

1. Installation of the AMX License Server application, which identifies existing licenses of AMX software products (including RMS Enterprise) present on the server. See Installing the AMX License Server below for details.
2. Entry of the RMS Entitlement information required to install and activate your RMS Enterprise Server and Client Licenses.

RMS does not support clustering of the AMX License Servers. The first RMS server installed will get the installation of AMX License Server on it. Subsequent RMS server installs will need to point back to the IP / hostname of the server where the AMX License service resides.

In production, when the server that hosts AMX License Server goes down, the other servers will continue to run uninterrupted for 24 hours. RMS has a 24 hour grace period where it will continue to run without communication to the AMX License Server.

Remote Log Files Directory

RMS servers write their log files to a local folder and to a remote shared folder. Both of these are configurable through the RMS Server Configuration tool. The remote log directory is needed because log files can be downloaded from the RMS UI. Having log files from all the RMS servers aggregated into a shared folder allows for easy access without having to log into each server individually.

Remote Data Files Directory

RMS indexes assets, locations, and client gateways for a rich searching experience. This search index is stored in a remote shared folder. All of the RMS servers read the index from this single place.

It's Your World. Take Control.

UNC Paths

When specifying the remote shared log and data directories in the RMS Server Configuration tool, the UNC path format should be used instead of referring to mapped Windows drive letters. The reason being that when Windows boots and launches the services, of which RMS is one, the drive letters aren't mapped yet. For a service, UNC path formats are required for accessing remote shared folders.

Cluster Names

RMS uses an underlying technology for clustering the servers together and passing information between them that requires special cluster names. The RMS Installer suggests default names that will work just fine for a single multi-server installation. However, if PV needs to configure more than one multi-server installation, then these cluster names need to be changed to be unique per multi-server installation. Otherwise, the two RMS installations would incorrectly find each other and participate as one even though they have different databases.

Worker Name

Ideally the worker name would match whatever is used for the Apache route name and Tomcat `jvmRoute` name, but it doesn't have to. The worker name setting on the Multi-Server Configuration page of the RMS Server Configuration tool really only controls the prefix to the server-specific log file that gets written to the Remote Log File Directory. So for example you would see the following log files in our shared directory:

```
Win2k8FrmSrv1_Server.log
Win2k8FrmSrv2_Server.log
Win2k8FrmSrv3_Server.log
Win2k8FrmSrv4_Server.log
```

It's Your World. Take Control.

Final Remarks

- Since the RMS application will be accessed through IIS, remember to not include port 8080 in the URL to the application (e.g. http://myserver.com/rms).
- Any clients that were previously configured to connect on port 8080 will need to be reconfigured to connect using port 80.
- The administrator needs to take software and network firewalls into consideration and make any necessary changes for this new deployment. IIS needs to talk to Tomcat over port 8009. If IIS and Tomcat are hosted on different machines or VMs, this port must be open.
- For more information regarding RMS system requirements or other aspects of the RMS system installation, please read the RMS Enterprise Installation Guide at: http://www.amx.com//assets/manuals/RMSEnterprise.InstallationGuide.pdf

Network Bandwidth

Network Measurement Methodology

The network traffic captures were done using Wireshark version 1.6.7. Each data capture was performed according to the following steps:

1. Start capture.
2. Perform the action to measure.
3. Wait until there is no more traffic.
4. Stop capture.
5. Save the capture.
6. Report results.

Capture of the frames used an AMX RPM project and a RMS Enterprise Cloud server. The frame sizes are shown below and the captures are attached.

Events

The Event Scenario does not necessarily require human intervention. Events occur either by a device triggering a threshold, a timed function occurring, or a system being brought on-line manually.

AMX
3000 Research Drive
Richardson, TX 75082

tel 800.222.0193
fax 469.624.7153

www.amx.com

It's Your World. Take Control.

23 | P a g e

Use Case: NetLinx System Connection
- When the RMS Client enters the ONLINE state, it means that the RMS Client Gateway has successfully been registered with the RMS system, has been licensed and approved by a RMS administrator, and assigned to a location in the RMS system.
- SYSTEM_ONLINE - 5862 bytes consisting of the following frames

BYTE FRAME	673	1389	1389	805	803	803

Use Case: Asset Power Mode On/Off
- Set Power (ON|OFF)
- On/Off Event – 1014 bytes (1 frame)

Use Case: Source Usage Activated
- This function is used to 'Activate' a source by its index number.
- RMS Source Usage Activate Source (Computer-1) – 990 bytes (1 frame)

Use Case: Lamp Consumption Update
- Parameter type LAMP_USAGE is updated
- Lamp Time – 995 bytes (1 frame)

Use Case: RMS Communications Ping
- RMS Server Ping and Ping Response
- Hearbeat – 803 bytes (1 frame)

Use Case: Request Help
- RMS can expose RMS-specific user interface functionality including help & maintenance requests and a location Hotlist view.
- RmsSendHelpRequest Help Request ("I need some help") – 970 bytes (1 frame)

It's Your World. Take Control.

RMS Enterprise Server Listener Sockets/Ports:

The RMS Enterprise server requires the following listener sockets and ports for inbound communication to the RMS server:

Listener Name/Protocol	Transport Protocol	Port	User Configurable	Notes
HTTP	TCP	80	YES	Used for endpoint device communication and web user interface. (Security Note*: NetLinx master endpoint device communication must communicate via HTTP (80). This port could be changed to an alternate port for obscurity, but it must use the HTTP protocol. One mitigation option would be to design a protected communication VLAN between the RMS server(s) and endpoints that is firewalled from public or general network access.*)
HTTPS	TCP	443	YES	Can optionally be used for web user interface with the installation and configuration of a SSL certificate. (Security Note: *NetLinx master endpoint device communication must communicate via port 80. Web user interface communication can be restricted to only use HTTPS 443*)
RMS v3.x (legacy)	TCP	3839	YES	Only required if the legacy option is enabled to support prior generation RMS 3.x client endpoints.

It's Your World. Take Control.

Listener Name/Protocol	Transport Protocol	Port	User Configurable	Notes
Proxy AJP	TCP	8009	YES	This is only used in deployments where another web server such as Apache Web Server or Microsoft IIS is sitting in front of the RMS server and proxying HTTP/HTTPS communication to the Tomcat web server hosting the RMS application AND the AJP proxy protocol is used. This is typically found in configurations where pre-authentication or Windows Integrated Authentication is configured for use. (Security Note: *If an AJP proxy is used ports 80/443 can be disabled or firewalled on the local RMS server machine; however they will need to be exposed on the proxy server.*)
SNMP	UDP	161	YES	Only required if implementing the optional integration features with an existing SNMP infrastructure.
Intra-server communication	UDP	45564	YES	Only required if implementing multiple RMS servers in a server cluster/failover configuration.
Intra-server communication	UDP	45588	NO	Only required if implementing multiple RMS servers in a server cluster/failover configuration.
Intra-server communication	UDP	46655	NO	Only required if implementing multiple RMS servers in a server cluster/failover configuration.

It's Your World. Take Control.

RMS Enterprise Server Outbound Sockets/Ports:

The RMS Enterprise server requires the following outbound sockets and ports for outbound communication from the RMS server:

Name/Protocol	Transport Protocol	Port	User Configurable	Notes
SMTP	TCP	25	YES	Only required if sending unsecured email communication to an unsecured SMTP server. (Security Note: *Use a secure SMTP server to mitigate this security risk. Alternatively, design a protected communication path between servers that is firewalled from public or general network access.*)
SMTP w/ TLS	TCP	587	YES	Only required if sending secure email communication to an SMTP server that supports TLS.
SMTP w/ SSL	TCP	465	YES	Only required if sending secure email communication to an SMTP server that supports SSL.
SNMP traps	UDP	162	YES	Only required if configured to send SNMP trap notification to an SNMP server/console.
LDAP	TCP	389	YES	Only required if implementing the optional integration features with an existing LDAP directory infrastructure for user authentication.
LDAPS	TCP	636	YES	Only required if implementing LDAP over SSL with a third-party certification authority.

It's Your World. Take Control.

Name/Protocol	Transport Protocol	Port	User Configurable	Notes
RMS Licensing Server	TCP	5093	NO	Only required if implementing multiple RMS servers in a server cluster/failover configuration OR if configured to use a remote Sentinel licensing server over the network. (Security Note: *If only deploying a single RMS server, install the licensing service on the same server to mitigate this security concern. Alternatively, design a protected communication path between servers that is firewalled from public or general network access.*)
Microsoft SQL Server	TCP	1433	YES	Only required if implementing multiple RMS servers in a server cluster/failover configuration OR if configured to use a remote MS SQL server/cluster over the network. (Security Note: *If only deploying a single RMS server and only supporting a small number of locations in the mitigate this security concern by installing the SQL server on the same machine. Alternatively, design a protected communication path between servers that is firewalled from public or general network access.*)
TRIBES (Intra-server communication)	UDP	45564	YES	Only required if implementing multiple RMS servers in a server cluster/failover configuration.
Hibernate Search + JGroups	UDP	45588	NO	Only required if implementing multiple RMS servers in a server cluster/failover configuration.
Infinispan + Jgroups	UDP	46655	NO	Only required if implementing multiple RMS servers in a server cluster/failover configuration.

It's Your World. Take Control.

Network infrastructure requirements

In a clustered deployment, the nodes must find each other using multicast UDP communication and all nodes must be on the same subnet.

By default, RMS uses the following multicast ports for a clustered deployment: 45564, 45588, and 46655.

Communication between the nodes provides a variety of functions:

1. Configuration changes made via the Flex UI. For example, when a configuration change within the web UI is made (e.g. SMTP Server), that change is persisted to the database and propagated to the other servers.

2. Search index write operations: when locations are created and client gateways and their assets are registered with the server, this information must be indexed on the shared file storage.

Only one server can write to the index at a time, so RMS employs a dynamic master-slave approach:

- The first server to start in a cluster becomes the master, and subsequent servers become slaves.
- The slaves send their index updates to the master who will perform the index writes for them.
- If the master node goes down, the second server that was started becomes the new master and performs the index writes.

3. Search index read operations (via SMB network file storage).

4. A distributed data grid (between the nodes) stores the last time each REST-based client gateway communicated with the server. This allows RMS to determine when a client has gone off-line.

5. Licensing checks. Each server must communicate with the AMX License Manager every 10 minutes.

If the AMX License Manager service goes down, RMS will continue to operate for a 24 hour grace period.

It's Your World. Take Control.

RMS Enterprise Client (Endpoint)

The RMS endpoint clients (*such as NetLinx masters/DVX/DGX*) server require the following client connection/consumer sockets and ports for network communication to the RMS server from the endpoint device.

| Name/Protocol	Transport Protocol	Port	User Configurable	Notes
HTTP	TCP	80	YES	Used for endpoint device communication and web user interface. (Security Note*: NetLinx master endpoint device communication must communicate via HTTP (80). This port could be changed to an alternate port for obscurity, but it must use the HTTP protocol. One mitigation option would be to design a protected communication VLAN between the RMS server(s) and endpoints that is firewalled from public or general network access.*)
HTTPS	TCP	443	YES	Can optionally be used for web user interface with the installation and configuration of a SSL certificate on the RMS server. (Security Note: *NetLinx master endpoint device communication must communicate via port 80. Web user interface communication can be restricted to only use HTTPS 443*)
RMS v3.x (legacy)	TCP	3839	YES	Only required if the legacy option is enabled to support prior generation RMS 3.x client endpoints.

It's Your World. Take Control.

RMS Enterprise Scheduling Interfaces

The RMS Scheduling Interface communicates with the RMS Server over HTTP. HTTPS is not supported for this interface. Both HTTP and HTTPS are supported for the connection between the RMS Interface and Scheduling system. To ensure optimal performance of the RMS Enterprise UI, the RMS Scheduling Interface application should not be installed on the Primary RMS Enterprise Server. Install the RMS Scheduling Interface application on a separate server. It is necessary to map each of the selected resources (Locations) in the RMS Enterprise Scheduling Configuration tool to a Resource Profile, in order to enable the scheduling interface for each location. This requires accessing the Location Management page in the RMS Enterprise UI. The RMS Enterprise UI is accessed via web browser.
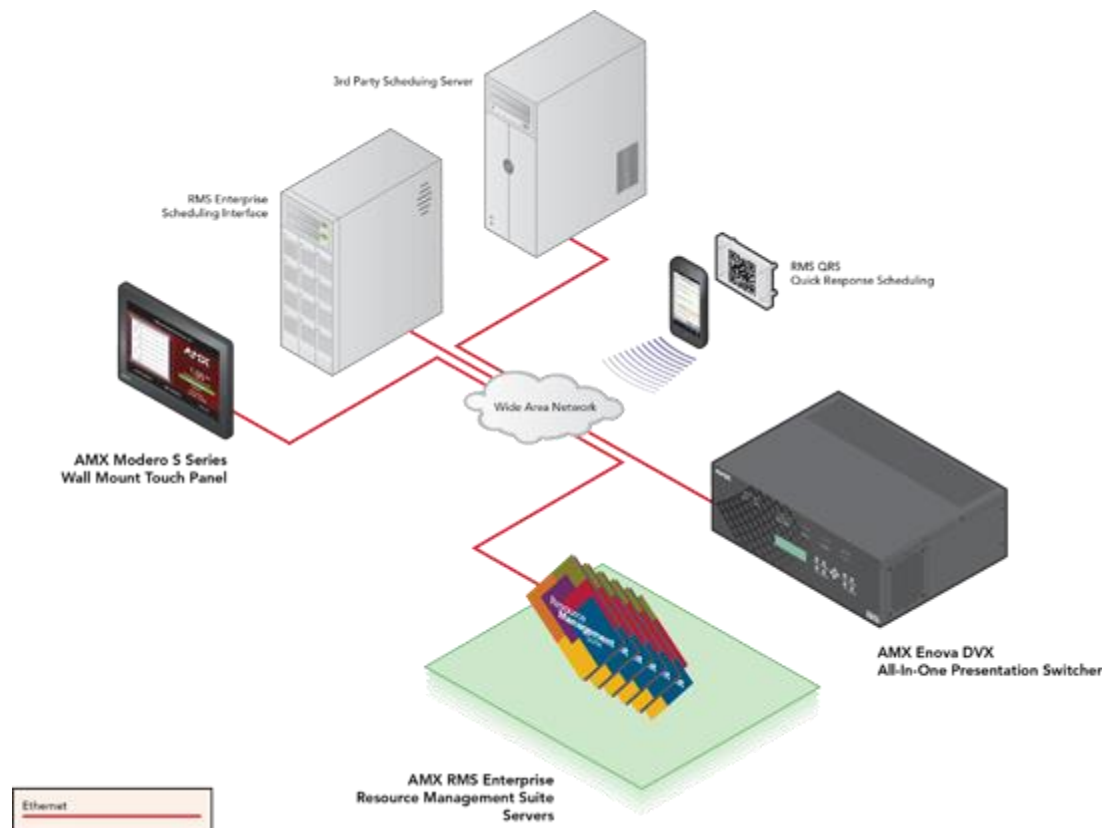
It's Your World. Take Control.

The RMS Scheduling Interface processes one location at a time when synchronizing appointments. This is not currently a parallelized operation. By default, the RMS Scheduling Interface pauses for 15 minutes between synchronization cycles. For example, if a synchronization takes 10 minutes to run for all locations, it will be idle for 15 minutes before the next troll.

The RMS Scheduling Interface requires the following client connection/consumer sockets and ports for network communication from the RMS scheduling service to the scheduling plug-in web services.

Name/Protocol	Transport Protocol	Port	User Configurable	Notes
HTTP	TCP	80	YES	Used for communication to a third party scheduling server/web services. (Security Note: *Use a secure HTTPS connection to Exchange to mitigate this security risk.*)
HTTPS	TCP	443	YES	Used for secure communication to a third party scheduling server/web services.

It's Your World. Take Control.

RMS Enterprise Database

Database Considerations for Maximum Performance and Scalability

Note that the overall performance of RMS Enterprise is a result of the server hardware and operating system used, as well as its configuration. Other factors include the number of Locations, Assets and Users in the system, as well as how the system is used. For example, the "Express Editions" of Microsoft SQL Server are appropriate to use for RMS Enterprise systems with less than 50 locations.

While this is true in most cases, it is important to note that an installation with a small number of rooms could be configured in such a way that it will generate a large amount of traffic to and from the server. As an example, a system with 50 locations, each of which contains a large number of devices with many control and monitoring functions running constantly, would certainly require at least the Standard version (possibly the Enterprise version) of Microsoft SQL Server in order to perform adequately, due to the large amount of traffic that would be generated.

When considering the server hardware to use with RMS Enterprise, it is important to understand not only the current requirements of the installation, but to also account for any potential upwards scaling of the system in the future.

It's Your World. Take Control.

Standard Practice for RMS Client Scalability

Typical conference and classrooms vary widely but certain device types and behavior may be assumed for the purpose of scalability evaluation. Typical polling rates for client to device is 30 seconds or greater and only changes to parameters on edge devices should be reported to the RMS server.

For the purpose of evaluation, a meeting room with two sources, video projector, screen, lighting system, and AV switcher was selected. The evaluation system utilized a single RMS server with separate SQL server and database meeting minimum required specifications. The evaluation consisted of 3000 client locations operating with similar behavior defined below.

Types of equipment used for data capture

- Lighting Control
- Screen Control
- Projector Control
- DVD Control
- Volume Control

Activities associated with described devices during a 30 minute session:

The system evaluation experienced heavy loads as all locations had meetings beginning on the hour and half hour with consider activity at simultaneous periods. This test was designed to generate as much network and processor traffic load as possible In order to ensure that minimum requirements met maximum uptime for the RMS system.

@ Start Time
1. Turn Light On
2. Light level 100%
3. Lower Screen
4. Turn Projector On
5. Set Source Input to 2
6. Set Volume
7. Turn DVD On
8. Set Lights to 50%

T= 15 Minutes
- Pause DVD

T= 17 Minutes
- Play DVD

T= 28 Minutes
1. Stop DVD
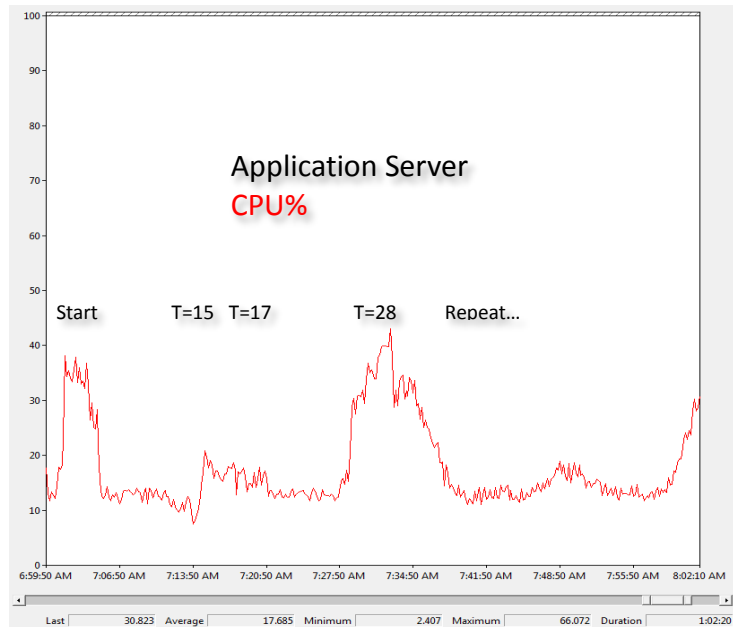2. Lights to 100%
3. DVD Power OFF
4. Projector Power OFF
5. Raise Screen
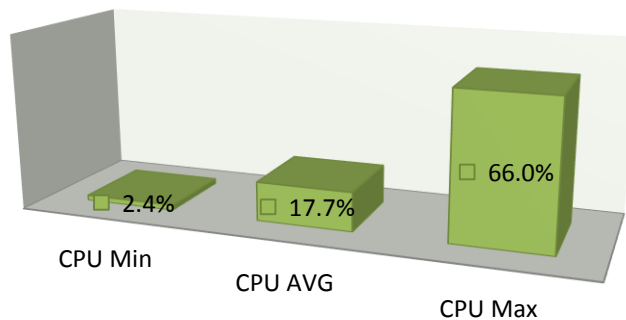6. Lights OFF

It's Your World. Take Control.

Evaluation system behavior is displayed in graph to the right and chart below. Note the processor performance peaks at 66% during high demand from 3000 locations simultaneously. Average CPU demand at 17% should be acceptable under average RMS client loading conditions.

Application Server
CPU%

| Start | T=15 | T=17 | T=28 | Repeat... |

| Last | 30.823 | Average | 17.685 | Minimum | 2.407 | Maximum | 66.072 | Duration | 1:02:20 |

Application Server

■ Application Server

2.4% CPU Min
17.7% CPU AVG
66.0% CPU Max

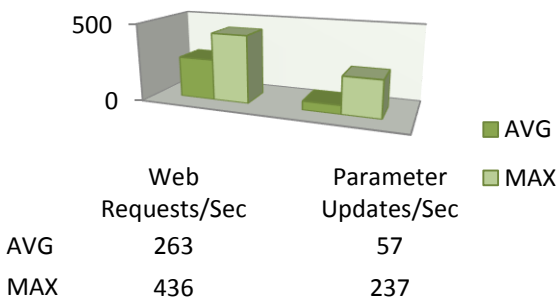It's Your World. Take Control.

Load limits

A web server (program) has defined load limits, because it can handle only a limited number of concurrent client connections (usually between 2 and 80,000, by default between 500 and 1,000) per IP address (and TCP port) and it can serve only a certain maximum number of requests per second depending on:

- Its own settings

- The HTTP request type

- Content origin (static or dynamic)

- The fact that the served content is or is not cached

- The hardware and software limitations of the OS where it is working

When a web server is near to or over its limits, it becomes unresponsive. Many factors influence the bandwidth costs, including the number of servers in the cluster as well as the number of locations, client gateways, and assets.

Running Room Scenario at normal time results in ~204,000 Parameter updates in an hour.

	Web Requests/Sec	Parameter Updates/Sec
AVG	263	57
MAX	436	237

It's Your World. Take Control.