Monitoring Microsoft Radius Server

eG Innovations Product Documentation

www.eginnovations.com

eG
Total Performance Visibility

About eG Innovations

eG Innovations provides intelligent performance management solutions that automate and dramatically accelerate the discovery, diagnosis, and resolution of IT performance issues in on-premises, cloud and hybrid environments. Where traditional monitoring tools often fail to provide insight into the performance drivers of business services and user experience, eG Innovations provides total performance visibility across every layer and every tier of the IT infrastructure that supports the business service chain. From desktops to applications, from servers to network and storage, from virtualization to cloud, eG Innovations helps companies proactively discover, instantly diagnose, and rapidly resolve even the most challenging performance and user experience issues.

eG Innovations is dedicated to helping businesses across the globe transform IT service delivery into a competitive advantage and a center for productivity, growth and profit. Many of the world's largest businesses use eG Enterprise to enhance IT service performance, increase operational efficiency, ensure IT effectiveness and deliver on the ROI promise of transformational IT investments across physical, virtual and cloud environments.

To learn more visit www.eginnovations.com.

Contact Us

For support queries, email support@eginnovations.com.

To contact eG Innovations sales team, email sales@eginnovations.com.

Table of Contents

Table of Figures

1

Chapter 1: Introduction

NPS (Network Policy Server) is the Microsoft implementation of a Remote Authentication Dial-in User Service (RADIUS) server, and as such, it performs connection authentication, authorization, and accounting for many types of network access, including wireless and virtual private network (VPN) connections. NPS also functions as a health evaluation server for NAP (Network Access Protection).

The following illustration shows NPS as a RADIUS server for a variety of access clients and a RADIUS proxy. NPS uses an Active Directory domain for user credential authentication of incoming RADIUS Access-Request messages.

When NPS is used as a RADIUS server, RADIUS messages provide authentication, authorization, and accounting for network access connections in the following way:

- Access servers, such as dial-up network access servers, VPN servers, and wireless access points, receive connection requests from access clients.
- The access server, configured to use RADIUS as the authentication, authorization, and accounting protocol, creates an Access-Request message and sends it to the NPS server.
- The NPS server evaluates the Access-Request message.
- If required, the NPS server sends an Access-Challenge message to the access server. The access server processes the challenge and sends an updated Access-Request to the NPS server.
- The user credentials are checked and the dial-in properties of the user account are obtained by using a secure connection to a domain controller.
- The connection attempt is authorized with both the dial-in properties of the user account and remote access policies.
- If the connection attempt is both authenticated and authorized, the NPS server sends an Access-Accept message to the access server. If the connection attempt is either not authenticated or not authorized, the NPS server sends an Access-Reject message to the access server.
- The access server completes the connection process with the access client and sends an Accounting-Request message to the NPS server, where the message is logged.
- The NPS server sends an Accounting-Response to the access server.

Issues in the functioning of NPS, if not promptly isolated and resolved, might result in the complete collapse of the remote authentication and authorization service provided by the Windows server. 24x7 monitoring of NPS, hence becomes imperative. The eG Enterprise Suite helps administrators in this task.

This document describes the monitoring model that eG Enterprise prescribes for Microsoft Radius server, and the performance metrics each model collects.

Chapter 2: Administering eG Manager to Work with Microsoft Radius server

To do the above, do the following:

1. Log into the eG administrative interface.
2. The Microsoft Radius server cannot be discovered by the eG Enterprise system. Therefore, proceed to add it using the **COMPONENTS** page that appears when the menu sequence, Infrastructure -> Components -> Add/Modify is followed. Remember that components manually added are managed automatically. Discovered components, however, are managed using the **COMPONENTS - MANAGE/UNMANAGE** page. Figure 2.1 clearly illustrates the process of adding an Microsoft Radius server.

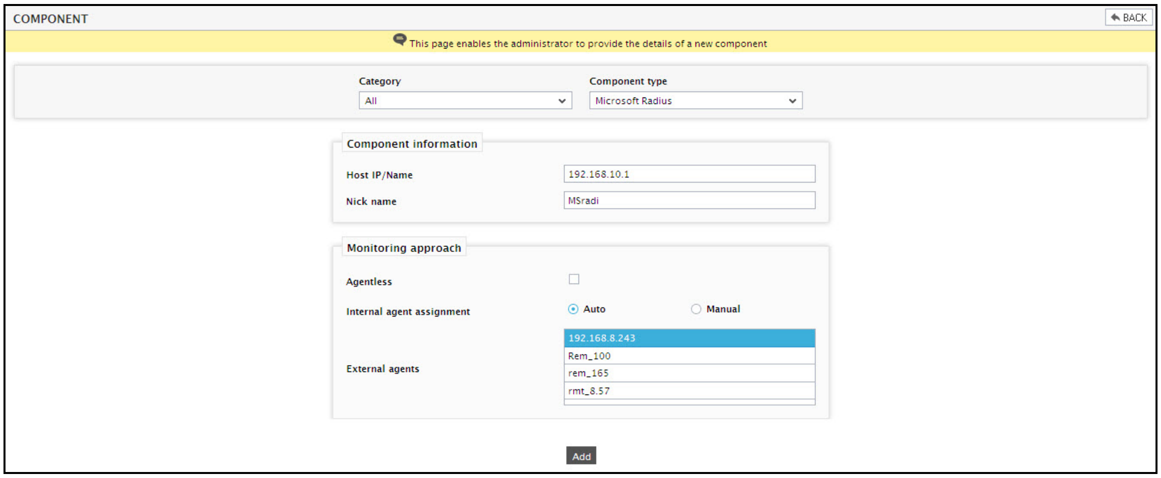

This page enables the administrator to provide the details of a new component

Category: All

Component type: Microsoft Radius

Component information

Host IP/Name: 192.168.10.1

Nick name: MSradi

Monitoring approach

Agentless: □

Internal agent assignment: ● Auto ○ Manual

External agents:
- 192.168.8.243
- Rem_100
- rem_165
- rmt_8.57

Figure 2.1: Adding a Microsoft Radius server

3. Specify the **Host IP** and **Nick name** for the Microsoft Radius server and click the **Add** button in the Figure 2.1 to register the changes.
4. Finally, sign out of the eG administrative interface.

3

Chapter 3: Monitoring Microsoft Radius Servers

The eG Enterprise suite provides out-of-the-box monitoring support to the Windows Internet Authentication Service, and proactively alerts administrators of authentication, authorization, or accounting bottlenecks encountered by the NPS server. The specialized *Microsoft Radius* monitoring model (see Figure 3.1) offered by the eG Enterprise suite executes a variety of tests on the NPS server; these tests, in turn, use the perfmon utility of Windows to extract critical performance statistics pertaining to the services offered by the NPS server.

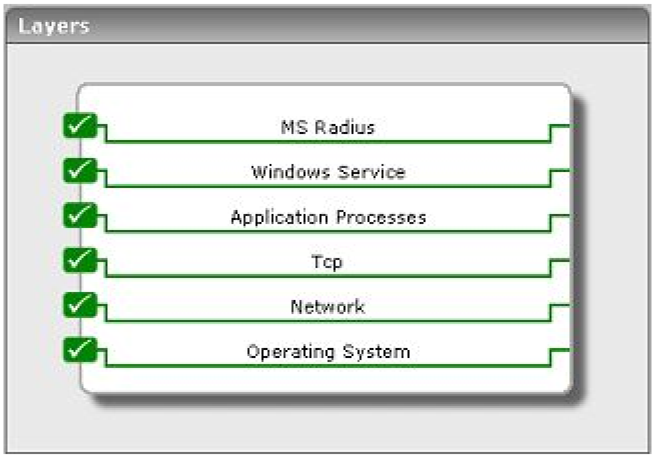

Figure 3.1: The layer model of the Microsoft Radius server

This section will discuss the **MS Radius** layer alone, as all the other layers have been elaborately discussed in the *Monitoring Unix and Windows Servers* document.

3.1 The MS Radius Layer

This layer monitors the authentication, authorization, and accounting activities performed by the NPS server and clients.

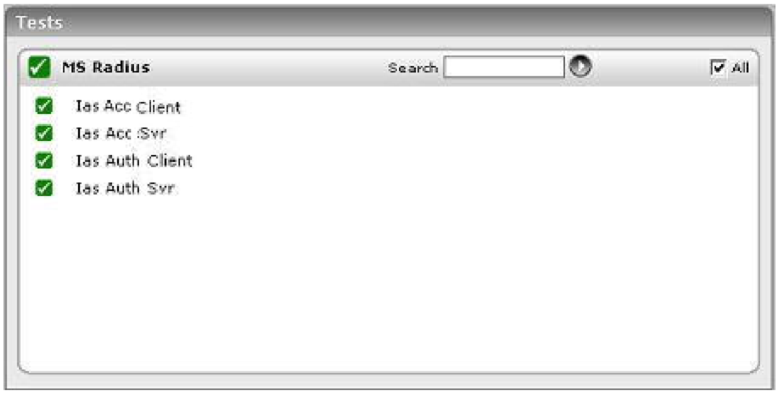

☑ Ias Acc Client
☑ Ias Acc Svr
☑ Ias Auth Client
☑ Ias Auth Svr

Figure 3.2: The tests associated with the MS Radius layer

3.1.1 NPS Accounting Server Test

Besides providing remote authentication services to clients, the NPS also provides a central accounting recording service for all accounting requests that are sent by the clients. Once the NPS server completes the connection process initiated by a client, the access server which processed the connection request sends an Accounting-Request message to the NPS server, where the message is logged. The NPS then sends an Accounting-Response to the access server. In addition, the access server also sends Accounting-Request messages for the following:

- During the time in which the connection is established
- When the access client connection is closed
- When the access server is started and stopped

This test monitors the accounting-requests received and accounting-responses sent by the NPS to clients.

Target of the test : An NPS server

Agent deploying the test : An internal agent

Outputs of the test : One set of results for every NPS server that is being monitored

Configurable parameters for the test

Parameters	Description
Test period	This indicates how often should the test be executed.
Host	The host for which the test is to be configured.
Port	The port at which the NPS server listens. The default is NULL.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Packets sent	Indicates the rate at which packets were sent by the NPS server.	Packets/Sec	
Packets received	Indicates the rate at which the NPS server received packets.	Packets/Sec	When viewed along with the Packets sent measure, this measure serves as a good indicator of the traffic on the server.
Packets dropped	Indicates the rate at which incoming packets were silently discarded for a reason other than being malformed, bad authenticators, or unknown types.	Packets/Sec	A consistent increase in the value of this measure is a cause for concern, and might warrant further investigation.
Invalid requests	Indicates the rate at which packets were received from an unknown address.	Reqs/Sec	
Malformed packets	Indicates the rate at which malformed packets were received; bad authenticators or unknown types are not included in this count.	Packets/Sec	
Unknown packets	Indicates the rate at which packets of an unknown type were received.	Packets/Sec	
No record packets	Indicates the rate at which RADIUS Accounting- Request	Records/Sec	

Measurement	Description	Measurement Unit	Interpretation
	packets were received and responded to but not recorded.		
Accounting requests	Indicates the rate at which RADIUS Accounting- Requests were received from this client on the accounting port.	Reqs/Sec	
Accounting responses	Indicates the rate at which RADIUS Accounting- Response packets were sent to this client on the accounting port.	Reqs/Sec	The Accounting requests and Accounting responses measures serve as effective indicators of the workload on the NPS server.
Duplicate requests	Indicates the rate at which duplicate RADIUS Accounting- Request packets were received from this client.	Reqs/Sec	
Bad authenticators	Indicates the rate at which Accounting-Requests containing invalid signature attributes were received.	Reqs/Sec	

3.1.2 NPS Accounting Client Test

This test monitors the accounting-requests sent and accounting-responses received by the RADIUS clients from the NPS servers.

Target of the test : An NPS server

Agent deploying the test : An internal agent

Outputs of the test : One set of results for every NPS server that is being monitored

Configurable parameters for the test

Parameters	Description
Test period	This indicates how often should the test be executed.
Host	The host for which the test is to be configured.
Port	The port at which the NPS server listens. The default is *NULL*.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Packets sent	Indicates the rate at which packets were sent by this client.	Packets/Sec	
Packets received	Indicates the rate at which this RADIUS client received packets.	Packets/Sec	When viewed along with the Packets sent measure, this measure serves as a good indicator of the traffic that originated from a client.
Packets dropped	Indicates the rate at which incoming packets were silently discarded for a reason other than being malformed, bad authenticators, or unknown types.	Packets/Sec	A consistent increase in the value of this measure is a cause for concern, and might warrant further investigation.
Malformed packets	Indicates the rate at which malformed packets were received; bad authenticators or unknown types are not included in this count.	Packets/Sec	
Unknown packets	Indicates the rate at which packets of an unknown type were received.	Packets/Sec	

Measurement	Description	Measurement Unit	Interpretation
No record packets	Indicates the rate at which RADIUS Accounting- Request packets were received and responded to but not recorded.	Records/Sec	
Accounting requests	Indicates the rate at which RADIUS Accounting- Requests were sent by this client on the accounting port.	Reqs/Sec	
Accounting responses	Indicates the rate at which RADIUS Accounting- Response packets were sent to this client on the accounting port.	Reqs/Sec	The Accounting requests and Accounting responses measures serve as effective indicators of the workload on the client.
Duplicate requests	Indicates the rate at which duplicate RADIUS Accounting- Request packets were received from this client.	Reqs/Sec	
Bad authenticators	Indicates the rate at which Accounting- Requests containing invalid signature attributes were received.	Reqs/Sec	

3.1.3 NPS Authentication Server Test

When NPS is used as a RADIUS server, it provides the a central authentication and authorization service for all access requests that are sent by RADIUS clients. NPS uses either a Microsoft® Windows NT® Server 4.0 domain, an Active Directory® domain, or the local Security Accounts

Manager (SAM) to authenticate user credentials for a connection attempt. NPS uses the dial-in properties of the user account and remote access policies to authorize a connection.

This test measures how well the NPS server performs remote authentication and authorization.

Target of the test : An NPS server

Agent deploying the test : An internal agent

Outputs of the test : One set of results for every NPS server that is being monitored

Configurable parameters for the test

Parameters	Description
Test period	This indicates how often should the test be executed.
Host	The host for which the test is to be configured.
Port	The port at which the NPS server listens. The default is NULL.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Packets sent	Indicates the rate at which packets were sent by the NPS server.	Packets/Sec	
Packets received	Indicates the rate at which the NPS server received packets.	Packets/Sec	When viewed along with the Packets sent measure, this measure serves as a good indicator of the traffic on the server.
Packets dropped	Indicates the rate at which incoming packets were silently discarded for a reason other than being malformed, bad authenticators, or unknown types.	Packets/Sec	A consistent increase in the value of this measure is a cause for concern, and might warrant further investigation.
Invalid requests	Indicates the rate at	Reqs/Sec	

Measurement	Description	Measurement Unit	Interpretation
	which packets were received from an unknown address.		
Malformed packets	Indicates the rate at which malformed packets were received; bad authenticators or unknown types are not included in this count.	Packets/Sec	
Unknown packets	Indicates the rate at which packets of an unknown type were received.	Packets/Sec	
Access accepts	Indicates the rate at which RADIUS Access-Accept packets were sent by the NPS server to this client.	Accepts/Sec	
Access challenges	Indicates the rate at which Access-Challenge messages are being processed.	Challenges/Sec	
Access rejects	Indicates the rate at which Access- Reject messages are being processed.	Rejects/Sec	A very high value of this measure could warrant a review of the remote access policies.
Access requests	Indicates the rate at which packets were received on an authentication port from this client.	Reqs/Sec	
Duplicate requests	Indicates the rate at which duplicate	Reqs/Sec	

Measurement	Description	Measurement Unit	Interpretation
	RADIUS Access-Request packets were received from this client.		

3.1.4 NPS Authentication Client Test

This test monitors the access-requests sent by access-clients to the NPS server, and indicates how many requests were accepted/rejected by the NPS server.

Target of the test : An NPS server

Agent deploying the test : An internal agent

Outputs of the test : One set of results for every NPS server that is being monitored

Configurable parameters for the test

Parameters	Description
Test period	This indicates how often should the test be executed.
Host	The host for which the test is to be configured.
Port	The port at which the NPS server listens. The default is NULL.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Packets sent	Indicates the rate at which packets were sent by this client to the NPS server.	Packets/Sec	
Packets received	Indicates the rate at which this client received packets from the NPS server.	Packets/Sec	When viewed along with the Packets sent measure, this measure serves as a good indicator of the traffic on the client.

Measurement	Description	Measurement Unit	Interpretation
Packets dropped	Indicates the rate at which incoming packets were silently discarded for a reason other than being malformed, bad authenticators, or unknown types.	Packets/Sec	A consistent increase in the value of this measure is a cause for concern, and might warrant further investigation.
Malformed packets	Indicates the rate at which malformed packets were received; bad authenticators or unknown types are not included in this count.	Packets/Sec	
Unknown packets	Indicates the rate at which packets of an unknown type were received.	Packets/Sec	
Access accepts	Indicates the rate at which RADIUS Access-Accept packets were sent to this client.	Accepts/Sec	
Access challenges	Indicates the rate at which Access-Challenge messages are being processed.	Challenges/Sec	
Access rejects	Indicates the rate at which Access- Reject messages are being processed.	Rejects/Sec	A very high value of this measure could warrant a review of the remote access policies.
Access requests	Indicates the rate at which packets were received on an authentication port from this client.	Reqs/Sec	

Measurement	Description	Measurement Unit	Interpretation
Bad authenticators	Indicates the rate at which packets containing invalid signature attributes were received.	Reqs/Sec	
Duplicate requests	Indicates the rate at which duplicate RADIUS Access-Request packets were received from this client.	Reqs/Sec	

3.1.5 NPS System Health Validators Test

NPS (Network Policy Server) is the Microsoft implementation of a Remote Authentication Dial-in User Service (RADIUS) server, and as such, it performs connection authentication, authorization, and accounting for many types of network access, including wireless and virtual private network (VPN) connections. NPS also functions as a health evaluation server for NAP (Network Access Protection).

System health validators (SHVs) in an NPS are server software counterparts to system health agents (SHAs) on NAP (Network Access Protection)-capable client computers. Each SHA on the client has a corresponding SHV in Network Policy Server (NPS). SHVs allow NPS to verify the statement of health (SoH) that is made by its corresponding SHA on the client computer. SHVs contain the details of the required configuration settings on client computers. For example, the Windows Security Health Validator (WSHV) is the counterpart to the Windows Security Health Agent (WSHA) on client computers. WSHV allows you to create a policy for the way in which settings on Network Access Protection (NAP)-capable client computers must be configured. If the settings on the client computer as reported in the SoH do not match the settings in the SHV on the server running NPS, it implies that the client computer is not compliant with the health policy requirements of the server. Once the system health validator validates the SoH from the client as either compliant or non-compliant, it marks the SoH with the relevant compliance status and sends it to the NPS.

By monitoring the statements of health issued by each system health validator, administrators can quickly capture non-compliances, investigate the reasons for the same, and can either fix it at the client side or fine-tune the access policies configured on the NPS to ensure secure access. This is exactly what the **NPS System Health Validators** test does. For each system health validator on NPS, this test reports the rate of compliances and non-compliances reported by that system health validator, thus shedding light on validations that often resulted in non-compliances. In addition, the test also reports the rate at which health statements could not be adjudged compliant/non-compliant, pinpoints the system health validators that sent out such statements, and reveals the reason for the same – is it owing to frequent server side failures? Or client side failures? Or is it because of other failures? The test also highlights 'slow' validators by measuring the responsiveness of every validator at pre-configured intervals.

Target of the test : An NPS server

Agent deploying the test : An internal agent

Outputs of the test : One set of results for every system health validator on the NPS server that is being monitored

Configurable parameters for the test

Parameters	Description
Test period	This indicates how often should the test be executed.
Host	The host for which the test is to be configured.
Port	The port at which the NPS server listens. The default is *NULL*.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Client-communication failures	Indicates the average number of Client-Communication failures per second from this health validator.	Failures/Sec	When a system health validator is not able to provide a health status to the Network Policy Server because of an error condition, it sends a Failure Category and code to the Network Policy Server. If the error is on the client side, the system health validator sends either a Client Component Failure Category or a Client Communication Failure Category.

Measurement	Description	Measurement Unit	Interpretation
			The value of these measures therefore indicate the rate at which client side failures occurred rendering a system health validator unable to determine the compliance status of a health statement. In the Configuration Manager System Health Validator properties on the Network Policy Server, errors tagged with these failure categories match to SHA not responding to NAP client and SHA unable to contact required services, respectively. Upon receipt of such failure categories from the system health validator, the NPS, by default, matches them to a non-compliant status.
Client- component failures	Indicates the average number of Client-Component failures per second from this health validator.	Failures/Sec	
Compliances	Indicates the rate at which compliant decsisions were issued to the NPS by this system health validator.	Decisions/Sec	This condition occurs when the client's compliant status is successfully validated by the System Health Validator point because all the following apply: • The statement of health is not older than the setting **Date created must be after**. • The statement of health is within the configured **Validity** period. • The client site is valid. • The client has used up-to-date

Measurement	Description	Measurement Unit	Interpretation
			Configuration Manager NAP policies. • A failure did not occur on either the Configuration Manager client or the System Health Validator point. A high value is desired for this measure.
Non-compliances	Indicates the rate at which non- compliant decsisions were issued to the NPS by this system health validator.	Decisions/Sec	This condition occurs when one of these situations apply: • The statement of health is older than the setting **Date created must be after**. • The statement of health is not within the configured **Validity** period. • The client does not have up-to-date Configuration Manager NAP policies. • The client has returned a non-compliant status because it does not have applicable software updates by the Effective Date as defined in the Configuration Manager NAP policies. A low value is desired for this measure.
None failures	Indicates the rate at which none failures were reported by this	Failures/Sec	

Measurement	Description	Measurement Unit	Interpretation
	system health validator.		
Other failures	Indicates the rate at which other failures were reported by this system health validator.	Failures/Sec	
Server-communication failures	Indicates the rate at which server-communication failures were reported by this system health validator.	Failures/Sec	When a system health validator is not able to provide a health status to the Network Policy Server because of an error condition, it sends a Failure Category and code to the Network Policy Server. If the error is on the server side, the system health validator sends either a Server Component Failure Category or a Server Communication Failure Category. The value of these measures therefore indicate the rate at which server side failures occurred rendering a system health validator unable to determine the compliance status of a health statement. In the Configuration Manager System Health Validator properties on the Network Policy Server, errors tagged with these failure categories match to SHV not responding and SHV unable to contact required services, respectively. Upon receipt of such failure categories from the system health validator, the NPS, by default,

Measurement	Description	Measurement Unit	Interpretation
			matches them to a non-compliant status.
Server-component failures	Indicates the rate at which server-component failures were reported by this system health validator.	Failures/Sec	
Last round-trip time	Indicates the interval (in hundredths of a second) between the most recent request to this system health validator and its response.	Secs	A low value is desired for this measure. A high value indicates that the system health validator is taking too long to validate health statements from the clients. Compare the value of this measure across system health validators to identify the slowest/least responsive validators.

3.1.6 NPS Remote Authentication Server Test

NPS performs centralized authentication, authorization, and accounting for wireless, authenticating switch, remote access dial-up and virtual private network (VPN) connections. When NPS is used as a RADIUS server, it provides a central authentication and authorization service for all access requests that are sent by RADIUS clients. NPS uses a Microsoft® Windows NT® Server 4.0 domain, an Active Directory® Domain Services (AD DS) domain, or the local Security Accounts Manager (SAM) user accounts database to authenticate user credentials for connection attempts.

The authenticating and authorization process is as follows:

1. Access servers, such as dial-up network access servers, VPN servers, and wireless access points, receive connection requests from access clients.
2. The access server, configured to use RADIUS as the authentication, authorization, and accounting protocol, creates an Access-Request message and sends it to the NPS server.
3. The NPS server evaluates the Access-Request message.
4. If required, the NPS server sends an Access-Challenge message to the access server. The access server processes the challenge and sends an updated Access-Request to the NPS server.

5. The user credentials are checked and the dial-in properties of the user account are obtained by using a secure connection to a domain controller.
6. The connection attempt is authorized with both the dial-in properties of the user account and network policies.
7. If the connection attempt is both authenticated and authorized, the NPS server sends an Access-Accept message to the access server.
8. If the connection attempt is either not authenticated or not authorized, the NPS server sends an Access-Reject message to the access server.

If NPS challenges access requests frequently or rejects requests very often, administrators need to be instantly notified of this, so that they can look into these aberrations and uncover their reasons. Likewise, administrators should also rapidly capture any unusual delay in request authentication by NPS, so that they can swiftly determine and fix the reason for the delay. For this, administrators should periodically run the **NPS Remote Authentication Server** test. This test tracks the Access-Request messages sent by every access server configured to use NPS for authentication, and reports the rate at which these access requests are challenged/rejected by NPS. In addition, the test reveals the time taken by NPS to authenticate requests to every server, thus proactively alerting administrators to potential slowdowns in authentication. The rate at which access requests to a server are enqueued on NPS pending processing is also revealed, so that administrators are informed of bottlenecks in authentication.

Target of the test : An NPS server

Agent deploying the test : An internal agent

Outputs of the test : One set of results for every access server that is configured to use NPS for authentication

Configurable parameters for the test

Parameters	Description
Test period	This indicates how often should the test be executed.
Host	The host for which the test is to be configured.
Port	The port at which the NPS server listens. The default is *NULL*.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Access-Accepts	Indicates the rate at which RADIUS Access-Accept packets were received by this server from NPS.	Accepts/Sec	This is a good indicator of how frequently access requests from clients to a server are authenticated and authorized by NPS.
Access-Challenges	Indicates the rate at which RADIUS Access-Challenge packets were sent by NPS to this server.	Challenges/Sec	A low value is desired for this measure. A high value indicates that NPS challenged many access requests, forcing the access server to send an updated Access-Request to NPS. In such cases, access clients are bound to experience delays in accessing the server.
Access-Rejects	Indicates the rate at which RADIUS Access-Reject packets were sent by NPS to this server.	Rejects/Sec	Ideally, the value of this measure should be 0 or very low. A high value indicates too many or too frequent request rejections, which in turn may cause access clients to be denied access to the server.
Access-Requests	Indicates the rate at which Access- Request packets were sent by this server to NPS.	Reqs/Sec	This is a good indicator of the load on NPS.
Bad authenticators	Indicates the rate at which this server sent access requests containing an invalid Message Authenticator attribute to NPS.	Reqs/Sec	Ideally, the value of this measure should be 0.

Measurement	Description	Measurement Unit	Interpretation
Packets dropped	Indicates the rate at which request packets sent by this server were silently discarded by NPS for a reason other than "malformed," "invalid Message Authenticator," or "unknown type".	Packets/Sec	Ideally, the value of this measure should be 0.
FullAccess-Decisions	Indicates the rate at which Full- access decisions were received from this server.	Decisions/Sec	NPS grants an access client full access if the client meets the defined health policies.
Malformed packets	Indicates the rate at which NPS received malformed packets from this server.	Packets/Sec	Ideally, the value of this measure should be 0.
Packets received	Indicates the rate at which requests packets were received from this server.	Packets/Sec	
Probation-Decisions	Indicates the rate at which probation-decisions were received from this server.	Decisions/Sec	If NPS grants an access client full access but for a limited period only, the client is said to be on probation. This can happen if NPS finds that the client did not fulfill certain health policy requirements.
Quarantine-Decisions	Indicates the rate at which quarantine decisions were sent by this server.	Decisions/Sec	When a remote access client dials in or connects via VPN to an access server, by default only the user's credentials (account name and password) are checked to

Measurement	Description	Measurement Unit	Interpretation
			determine whether access is granted. This means a computer that does not meet the network's policy requirements could still connect to the server and the network from a remote location. When quarantine control is deployed, after the user's credentials are authenticated the connection is "quarantined." In quarantine mode, the computer has an IP address and has limited access to some network resources (called quarantine resources) such as a DNS server and perhaps a file server or web server from which it can download files necessary to comply with the policies or where the user can get more information, but cannot access the rest of the network.
Request timeouts	Indicates the rate at which requests to this server timed out.	Reqs/Sec	A high value indicates frequent timeouts. Under such circumstances, you may want to consider changing the timeout setting for requests, so that timeouts are kept at a minimum.
Retransmissions	Indicates the rate at which requests were retransmitted to this server.	Reqs/Sec	Retransmits can increase the number of requests to NPS, thus overloading it. It is hence good practice to keep the rate of retransmissions minimal.

Measurement	Description	Measurement Unit	Interpretation
			One of the reasons for a high rate of retransmissions is a low Timeout setting on NPS. If the value of this measure is very high, you may want to change the timeout setting to reduce retransmits.
Unknown type	Indicates the average number of unknown type (non- RADIUS) packets received by this server per second.	Packets/Sec	
Last round-trip time	Indicates the interval (in hundredths of a second) between the most recent request to a remote NPS server and its response.	Secs	Ideally, the value of this measure should be very low. A high value indicates that that NPS is taking too long to authenticate requests.
Pending requests	Indicates the rate of requests destined for this server that have not yet timed out or received a response.	Reqs/Sec	A high value could either indicate a processing bottleneck on NPS or a high timeout setting. In the case of the latter, you may want to consider modifying the timeout setting to minimize the number of pending requests.

3.1.7 NPS Remote Accounting Server Test

Network Policy Server (NPS) supports Remote Authentication Dial-In User Service (RADIUS) accounting, which you can use to track network usage for auditing and billing purposes. Accounting data can also be queried to assist with network access troubleshooting.

When a RADIUS client is configured to use RADIUS accounting, at the start of service delivery it generates an Accounting-Start message describing the type of service being delivered and the user it is being delivered to. The message is then sent to the RADIUS Accounting server, which sends back an acknowledgment to the RADIUS client. At the end of service delivery, the client generates an Accounting-Stop message describing the type of service that was delivered and optional statistics, such as elapsed time, input and output octets, or input and output packets. It then sends that data to the RADIUS accounting server, which sends back an acknowledgment to the RADIUS client.

The Accounting-Request message (whether for the Start or Stop message) is submitted to the RADIUS accounting server through the network. If the quality of this network connection is poor, then many request packets may be malformed, forcing the server to drop them. This in turn can delay or deny responses to accounting servers and acknowledgments to clients.

Also, if the accounting server is overloaded with requests, the server can choke slowing down accounting in the bargain.

To avoid such delays, administrators must track accounting requests and responses, proactively detect potential slowdowns, accurately isolate what is causing it, and promptly fix it. This is where the **NPS Remote Accounting Server** test helps.

This test tracks the requests to and responses from each accounting server and reveals whether/not the servers are responding as quickly as the requests come in. You can also use this test to monitor the time each server takes to process requests, and thus identify the server that is experiencing a processing bottleneck. In addition, the test also captures the rate at which packets are dropped by the server and malformed/erroneous packets are received by the server, thus pointing to issues with the client or in the network connection between the server and the client. The load on each server is also revealed by monitoring the packets received by and the pending requests on the server from time to time. This way, the test pinpoints irregularities in load-balancing between servers, which in time can bottleneck request processing by the servers, resulting in slowdowns.

Target of the test : An NPS server

Agent deploying the test : An internal agent

Outputs of the test : One set of results for every access server that is configured to use NPS for authentication

Configurable parameters for the test

Parameters	Description
Test period	This indicates how often should the test be executed.
Host	The host for which the test is to be configured.

Parameters	Description
Port	The port at which the NPS server listens. The default is NULL.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Accounting-Requests	Indicates the rate at which this server receives accounting requests.	Reqs/Sec	This is a good indicator of the workload on the server. Compare the value of this measure across servers to know which server is overloaded. If a single server appears to be handling a vast majority of the requests, you may want to consider sprucing up your load-balancing algorithm, so that the request load is uniformly balanced across servers. If required, you can even consider adding more servers.
Accounting - Responses	Indicates the rate at which this server is responding to requests.	Reqs/Sec	If the value of this measure is much lower than the value of the Accounting-Requests measure, it could indicate that the server is not responding to requests quickly. You may want to investigate the reasons for the same.
Bad authenticators	Indicates the rate at which this server received requests containing an invalid Message Authenticator attribute.	Reqs/Sec	Ideally, the value of this measure should be 0.
Packets dropped	Indicates the rate at which this server were silently discarded the	Packets/Sec	Ideally, the value of this measure should be 0.

Measurement	Description	Measurement Unit	Interpretation
	request packets it received for a reason other than "malformed," "invalid Message Authenticator," or "unknown type".		
Malformed packets	Indicates the rate at which this server received malformed packets.	Packets/Sec	Ideally, the value of this measure should be 0.
Packets received	Indicates the rate at which requests packets were received by his server.	Packets/Sec	
Request timeouts	Indicates the rate at which requests to this server timed out.	Reqs/Sec	A high value indicates frequent timeouts. Under such circumstances, you may want to consider changing the timeout setting for requests, so that timeouts are kept at a minimum.
Retransmissions	Indicates the rate at which requests were retransmitted to this server.	Reqs/Sec	Retransmits can increase the number of requests to the server, thus overloading it. It is hence good practice to keep the rate of retransmissions minimal. One of the reasons for a high rate of retransmissions is a low Timeout setting on the server. If the value of this measure is very high, you may want to change the timeout setting to reduce retransmits.

Measurement	Description	Measurement Unit	Interpretation
Unknown type	Indicates the average number of unknown type (non- RADIUS) packets received by this server per second.	Packets/Sec	
Last round-trip time	Indicates the interval (in hundredths of a second) between the most recent request to this server and its response.	Secs	Ideally, the value of this measure should be very low. A high value indicates that that the accounting server is taking too long to perform accounting.
Pending requests	Indicates the rate of requests destined for this server that have not yet timed out or received a response.	Reqs/Sec	A high value could either indicate a processing bottleneck on the server or a high timeout setting (which could be causing many requests to be retransmitted to the server). In the case of the latter, you may want to consider modifying the timeout setting to minimize the number of pending requests.

3.1.8 NPS Policy Engine Test

Every network policy must have at least one configured condition. NPS provides many conditions groups that allow you to clearly define the properties that the connection request received by NPS must have in order to match the policy. How quickly NPS matches requests with policies is a good measure of the efficiency of the NPS policy engine. Using the **NPS Policy Engine** test, administrators can measure just that! This test reports the time taken by NPS to process requests, the rate of pending requests on NPS, and the number of requests that matched configured policies. In the process, the test reveals processing bottlenecks on the NPS and how they impact policy matching.

Target of the test : An NPS server

Agent deploying the test : An internal agent

Outputs of the test : One set of results for every NPS server that is being monitored

Configurable parameters for the test

Parameters	Description
Test period	This indicates how often should the test be executed.
Host	The host for which the test is to be configured.
Port	The port at which the NPS server listens. The default is NULL.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Last round-trip time	Indicates the interval (in hundredths of a second) between the most recent request to NPS and its response.	Secs	Ideally, the value of this measure should be very low. A high value indicates that NPS is taking too long to verify whether/not requests it receives match with the policy configuration.
Matched remote access policy	Indicates the average number of remote access policies that have been matched with requests per second.	Number	
Pending requests	Indicates the rate of requests destined for NPS that have not yet timed out or received a response.	Reqs/Sec	A high value could either indicate a processing bottleneck on NPS or a high timeout setting (which could be causing many requests to be retransmitted to the NPS). In the case of the latter, you may want to consider modifying the timeout setting to minimize the number of pending requests.

3.1.9 NPS Authentication Proxy Test

Network Policy Server (NPS) can be used as a RADIUS proxy to provide the routing of RADIUS messages between RADIUS clients access servers and RADIUS servers that perform user authentication, authorization, and accounting for the connection attempt. When used as a RADIUS proxy, NPS is a central switching or routing point through which RADIUS access and accounting messages flow.

The below shows NPS as a RADIUS proxy between RADIUS clients (access servers) and either RADIUS servers or another RADIUS proxy.

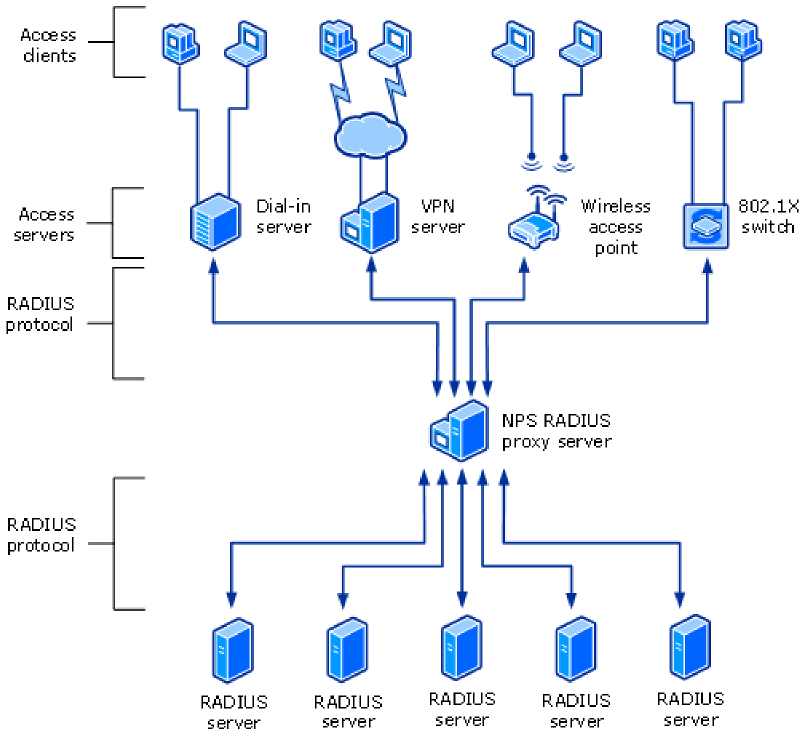

Figure 3.3: How NPS RADIUS Proxy works

When NPS is used as a RADIUS proxy between a RADIUS client and a RADIUS server, RADIUS messages for network access connection attempts are forwarded in the following way:

1. Access servers, such as dial-up network access servers, virtual private network (VPN) servers, and wireless access points, receive connection requests from access clients.
2. The access server, configured to use RADIUS as the authentication, authorization, and accounting protocol, creates an Access-Request message and sends it to the NPS server that is being used as the NPS RADIUS proxy.

3. The NPS RADIUS proxy receives the Access-Request message and, based on the locally configured connection request policies, determines where to forward the Access-Request message.
4. The NPS RADIUS proxy forwards the Access-Request message to the appropriate RADIUS server.
5. The RADIUS server evaluates the Access-Request message.
6. If required, the RADIUS server sends an Access-Challenge message to the NPS RADIUS proxy, where it is forwarded to the access server. The access server processes the challenge with the access client and sends an updated Access-Request to the NPS RADIUS proxy, where it is forwarded to the RADIUS server.
7. The RADIUS server authenticates and authorizes the connection attempt.
8. If the connection attempt is both authenticated and authorized, the RADIUS server sends an Access-Accept message to the NPS RADIUS proxy, where it is forwarded to the access server.
9. If the connection attempt is either not authenticated or not authorized, the RADIUS server sends an Access-Reject message to the NPS RADIUS proxy, where it is forwarded to the access server.

If the RADIUS server challenges or rejects connection requests frequently, the NPS RADIUS Proxy will transmit the same to the access servers. This in turn can cause many clients to be denied connections to access servers; some others may experience significant delays in connecting. Malformed request packets and those with invalid attributes/addresses can also be responsible for authentication delays/denials. Sometimes, a processing bottleneck on the NPS RADIUS Proxy server can also result in a slowdown in authentication. To enable clients to access remote services rapidly, administrators should keep a close watch on how the NPS RADIUS Proxy handles the requests and responses it receives, detect abnormalities rapidly, and quickly initiate measures to resolve them. The **NPS Authentication Proxy** test helps administrators do just that!

This test keeps tabs on the access requests the NPS RADIUS Proxy receives from access servers and reports the rate at which the RADIUS server challenges / rejects these requests, thereby enabling administrators to instantly spot an abnormal number of challenges and rejections. Additionally, the test also reveals the rate at which erroneous request/response packets are received by the NPS RADIUS Proxy, thus providing administrators with effective pointers to what could be causing the high rate of challenges, rejects, or retransmisssions – is it because of malformed packets? packets with invalid authentication attributes? packets with invalid addresses? or non-RADIUS packets? The test also sheds light on the poor processing ability of the NPS RADIUS Proxy by reporting the number of pending requests on the proxy from time to time.

Target of the test : An NPS server

Agent deploying the test : An internal agent

Outputs of the test : One set of results for the NPS RADIUS Proxy

Configurable parameters for the test

Parameters	Description
Test period	This indicates how often should the test be executed.
Host	The host for which the test is to be configured.
Port	The port at which the NPS server listens. The default is NULL.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Access-Accepts	Indicates the rate at which RADIUS Access-Accept packets were received by the proxy.	Accepts/Sec	
Access-Challenges	Indicates the rate at which RADIUS Access-Challenge packets were received by the NPS RADIUS proxy from the RADIUS server.	Challenges/Sec	A low value is desired for this measure. A high value indicates that the RADIUS server challenged many access requests, forcing the access server to send an updated Access- Request to it, through the NPS RADIUS proxy. In such cases, access clients are bound to experience delays in accessing the server.
Access-Rejects	Indicates the rate at which RADIUS Access-Reject packets were sent by the RADIUS server to the NPS	Rejects/Sec	Ideally, the value of this measure should be 0 or very low. A high value indicates too many or too frequent request rejections, which in turn may

Measurement	Description	Measurement Unit	Interpretation
	RADIUS Proxy.		cause access clients to be denied access to the server.
Access-Requests	Indicates the rate at which Access- Request packets were received by the NPS RADIUS Proxy from the access servers.	Reqs/Sec	This is a good indicator of the load on NPS.
Bad authenticators	Indicates the rate at which NPS RADIUS Proxy received access requests containing an invalid Message Authenticator attribute.	Reqs/Sec	Ideally, the value of this measure should be 0.
Packets dropped	Indicates the rate at which request packets received by the NPS RADIUS proxy were silently discarded for a reason other than "malformed," "invalid Message Authenticator," or "unknown type".	Packets/Sec	Ideally, the value of this measure should be 0.
FullAccess-Decisions	Indicates the rate at which Full- access decisions were received the NPS RADIUS Proxy.	Decisions/Sec	The RADIUS server grants an access client full access if the client meets the defined health policies.
Invalid addresses	Indicates the rate at which the NPS RADIUS Proxy received packets from unknown addresses.	Packets/Sec	Ideally, this value should be 0.

Measurement	Description	Measurement Unit	Interpretation
Malformed packets	Indicates the rate at which the NPS RADIUS Proxy received malformed packets.	Packets/Sec	Ideally, the value of this measure should be 0.
Packets received	Indicates the rate at which packets were received by the NPS RADIUS Proxy.	Packets/Sec	
Probation-Decisions	Indicates the rate at which probation-decisions were received from the NPS RADIUS Proxy.	Decisions/Sec	If the RADIUS server grants an access client full access but for a limited period only, the client is said to be on probation. This can happen if the RADIUS server finds that the client did not fulfill certain health policy requirements.
Quarantine-Decisions	Indicates the rate at which quarantine decisions were sent from the NPS RADIUS Proxy.	Decisions/Sec	When a remote access client dials in or connects via VPN to an access server, by default only the user's credentials (account name and password) are checked to determine whether access is granted. This means a computer that does not meet the network's policy requirements could still connect to the server and the network from a remote location. When quarantine control is deployed, after the user's credentials are authenticated the connection is "quarantined." In quarantine mode, the computer has an IP address and has

Measurement	Description	Measurement Unit	Interpretation
			limited access to some network resources (called quarantine resources) such as a DNS server and perhaps a file server or web server from which it can download files necessary to comply with the policies or where the user can get more information, but cannot access the rest of the network.
Request timeouts	Indicates the rate at which requests to the NPS RADIUS proxy timed out.	Reqs/Sec	A high value indicates frequent timeouts. Under such circumstances, you may want to consider changing the timeout setting for requests, so that timeouts are kept at a minimum.
Retransmissions	Indicates the rate at which requests were retransmitted to the NPS RADIUS Proxy.	Reqs/Sec	Retransmits can increase the number of requests to the proxy, thus overloading it. It is hence good practice to keep the rate of retransmissions minimal. One of the reasons for a high rate of retransmissions is a low Timeout setting on NPS RADIUS Proxy. If the value of this measure is very high, you may want to change the timeout setting to reduce retransmits.
Unknown type	Indicates the average number of unknown type (non- RADIUS) packets received by this	Packets/Sec	

Measurement	Description	Measurement Unit	Interpretation
	servethe NPS RADIUS proxy per second.		
Pending requests	Indicates the rate of requests destined for the proxy that have not yet timed out or received a response.	Reqs/Sec	A high value could either indicate a processing bottleneck on the NPS RADIUS Proxy or a high timeout setting. In the case of the latter, you may want to consider modifying the timeout setting to minimize the number of pending requests. In the case of the former, you may want to consider adding more RADIUS servers, so that the NPS RADIUS Proxy is able to dynamically balance the load of connection requests across multiple RADIUS servers and thus speed up processing.

3.1.10 NPS Accounting Proxy Test

Once the access server completes the connection process with the access client, it sends an Accounting-Request message to the NPS RADIUS proxy. The NPS RADIUS proxy logs the accounting data and forwards the message to the RADIUS server. The RADIUS server then sends an Accounting-Response to the NPS RADIUS proxy, where it is forwarded to the access server.

If the access servers experience accounting delays, it could either be owing to a slowdown in the NPS RADIUS proxy that routes the responses or the poor processing ability of the RADIUS server that sends the responses to the proxy. Malformed packets, packets with invalid attributes/addresses, and non-RADIUS packets can aso contribute to the time lag at the proxy server end. Another common reason for the slowdown is a request overload on the NPS RADIUS proxy.

Using the **NPS Accounting Proxy** test, administrators can measure how effectively the proxy is handling accounting requests, detect slowdowns, and pinpoint the probable reasons for the same.

Target of the test : An NPS server

Agent deploying the test : An internal agent

Outputs of the test : One set of results for the NPS RADIUS Proxy

Configurable parameters for the test

Parameters	Description
Test period	This indicates how often should the test be executed.
Host	The host for which the test is to be configured.
Port	The port at which the NPS server listens. The default is NULL.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Accounting-Requests	Indicates the rate at which the NPS RADIUS proxy receives accounting requests.	Reqs/Sec	This is a good indicator of the workload on the server.
Accounting - Responses	Indicates the rate at which the NPS RADIUS proxy is responding to requests.	Reqs/Sec	If the value of this measure is much lower than the value of the Accounting-Requests measure, it could indicate that the server is not responding to requests quickly. You may want to investigate the reasons for the same. This could either be caused by a processing bottleneck on the proxy server, a poor network connection between the proxy and the RADIUS server, or a slowdown on the RADIUS server.
Bad authenticators	Indicates the rate at which the proxy received requests containing an invalid Message Authenticator attribute.	Reqs/Sec	Ideally, the value of this measure should be 0.
Packets dropped	Indicates the rate at	Packets/Sec	Ideally, the value of this measure

Measurement	Description	Measurement Unit	Interpretation
	which the proxy silently discarded the request packets it received for a reason other than "malformed," "invalid Message Authenticator," or "unknown type".		should be 0.
Invalid addresses	Indicates the rate at which the proxy received packets with invalid addresses.	Reqs/Sec	Ideally, the value of this measure should be 0.
Malformed packets	Indicates the rate at which the proxy received malformed packets.	Packets/Sec	Ideally, the value of this measure should be 0.
Packets received	Indicates the rate at which requests packets were received by the proxy.	Packets/Sec	
Request timeouts	Indicates the rate at which requests to the proxy timed out.	Reqs/Sec	A high value indicates frequent timeouts. Under such circumstances, you may want to consider changing the timeout setting for requests, so that timeouts are kept at a minimum.
Retransmissions	Indicates the rate at which requests were retransmitted to the proxy.	Reqs/Sec	Retransmits can increase the number of requests to the proxy server, thus overloading it. It is hence good practice to keep the rate of retransmissions minimal. One of the reasons for a high rate of retransmissions is a low Timeout setting on the server.

Measurement	Description	Measurement Unit	Interpretation
			If the value of this measure is very high, you may want to change the timeout setting to reduce retransmits.
Unknown type	Indicates the average number of unknown type (non- RADIUS) packets received by the proxy per second.	Packets/Sec	
Pending requests	Indicates the rate of requests destined for the proxy that have not yet timed out or received a response.	Reqs/Sec	A high value could either indicate a processing bottleneck on the proxy, a high timeout setting on the proxy (which could be causing many requests to be retransmitted to the server), or the poor processing power of the RADIUS server. A flaky network connection between the proxy and the RADIUS server can also contribute to the processing delay and add to the count of pending requests.

4

Chapter 4: Conclusion

This document has described in detail the monitoring paradigm used and the measurement capabilities of the eG Enterprise suite of products with respect to **Microsoft Radius** server. For details of how to administer and use the eG Enterprise suite of products, refer to the user manuals.

We will be adding new measurement capabilities into the future versions of the eG Enterprise suite. If you can identify new capabilities that you would like us to incorporate in the eG Enterprise suite of products, please contact support@eginnovations.com . We look forward to your support and cooperation. Any feedback regarding this manual or any other aspects of the eG Enterprise suite can be forwarded to feedback@eginnovations.com.