Monitoring Double-Take Availability Server

eG Innovations Product Documentation

www.eginnovations.com

eG
Total Performance Visibility

Table of Contents

Table of Figures

Chapter 1: Introduction

Double-Take Availability for Windows provides real-time high availability and immediate disaster recovery so you never have to worry about downtime or the lost revenue and chaos that ensue.

Double-Take Availability ensures the availability of critical workloads. Using real-time replication and failover, you can protect data, individual applications, entire servers, or virtual machines. Identify your critical workload on your production server, known as the source, and replicate the workload to a backup server, known as the target. The target server, on a local network or at a remote site, stores the copy of the workload from the source. Double-Take Availability monitors any changes to the source workload and sends the changes to the copy stored on the target server. By replicating only the file changes rather than copying an entire file, Double-Take Availability server allows you to more efficiently use resources.

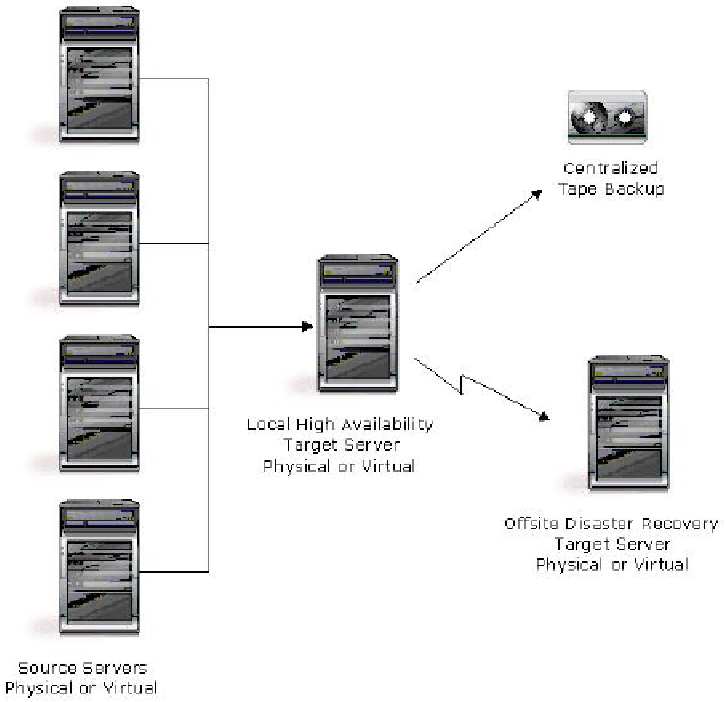

Figure 1.1: How does Double-Take Availability server work?

In physical/virtual infrastructures that deliver mission-critical services to end-users, the 24x7 availability of data and applications is crucial for maximizing service quality, user satisfaction, and consequently, revenues. In such environments therefore, the uninterrupted functioning of Double-Take Availability server is imperative. Issues such as intermittent breaks in the availability of Double-Take Availability server, excessive load conditions, and delayed connectivity, if not promptly detected and resolved, can significantly impact the delivery of underlying services. You thus need to periodically monitor Double-Take Availability server for such availability and operational snags, and initiate early measures to redress them. This is where the eG Enterprise lends helping hands to administrators for continuously monitoring the Double-Take Availability server.

Chapter 2: How to Monitor Double-Take Availability Server using eG Enterprise?

eG Enterprise monitors the Double-Take Availability Server in an agentless manner. All that is required for this is a single eG agent on any remote Windows host in the environment. This agent is capable of polling the SNMP MIB Of the server at regular intervals and fetching statistics related its performance. Before attempting to monitor the server, ensure that the Double-Take Availability server is SNMP enbled.

2.1 Managing the Double-Take Availability Server

The eG Enterprise cannot automatically discover the Double-Take Availability server so that you need to manually add the component for monitoring. Remember that the eG Enterprise automatically manages To manage a Double-Take Availability component, do the following:

1. Log into the eG administrative interface.
2. Follow the Components -> Add/Modify menu sequence in the **Infrastructure** tile of the **Admin** menu.
3. In the **COMPONENT** page that appears next, select Double-Take Availability as the **Component type**. Then, click the **Add New Component** button. This will invoke Chapter 2.

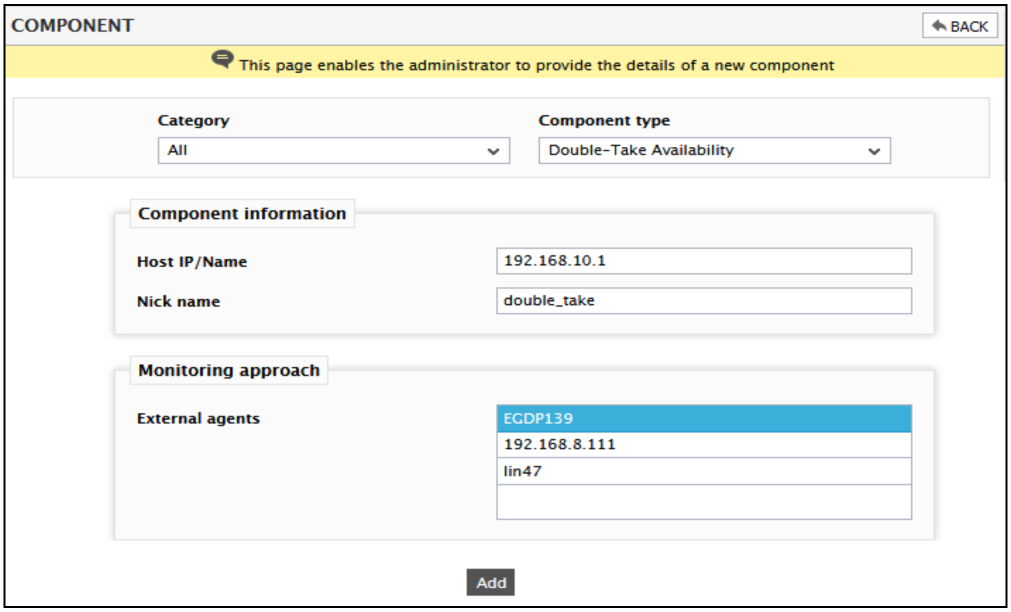

COMPONENT

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

Category: All

Component type: Double-Take Availability

Component information

Host IP/Name: 192.168.10.1

Nick name: double_take

Monitoring approach

External agents: EGDP139, 192.168.8.111, lin47

Figure 2.1: Adding the Double-Take Availability server

4. Specify the **Host IP/Name** and the **Nick name** of the Double-Take Availability server in Chapter 2. then, click on the **Add** button to register the changes.
5. Now, when you attempt to sign out of the eG administrative interface, Figure 2.2 appears, listing the tests requiring manual configuration.

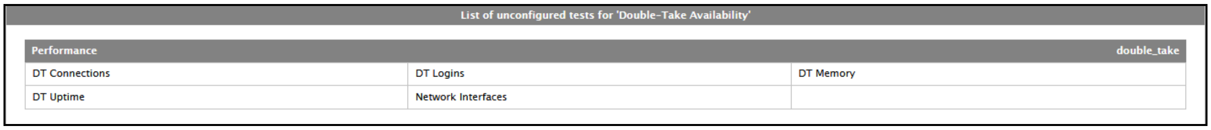

List of unconfigured tests for 'Double-Take Availability'

Performance		double_take
DT Connections	DT Logins	DT Memory
DT Uptime	Network Interfaces	

Figure 2.2: The list of Unconfigured tests for the DoubleTake server

6. Click on any test in the list of unconfigured tests. For instance, click on the **DT Logins** test to configure it. In the page that appears, specify the parameters as shown in Figure 2.2.

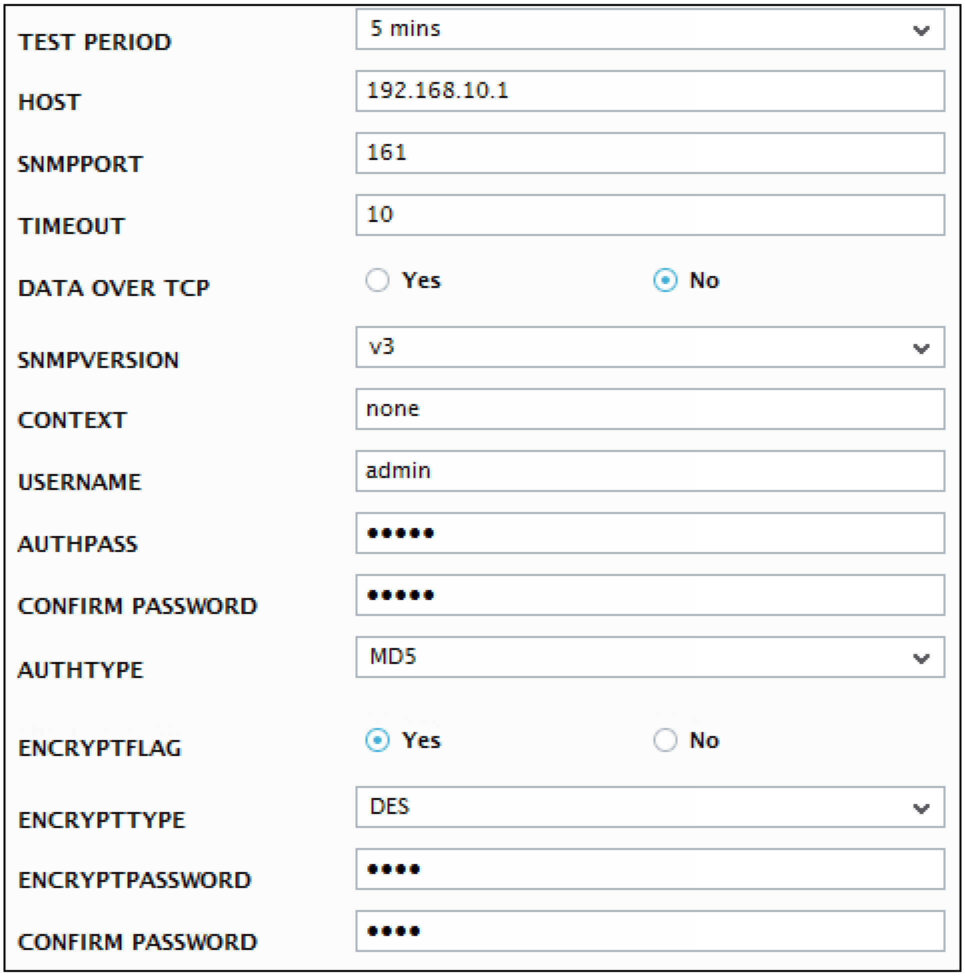

TEST PERIOD	5 mins	
HOST	192.168.10.1	
SNMPPORT	161	
TIMEOUT	10	
DATA OVER TCP	○ Yes	● No
SNMPVERSION	v3	
CONTEXT	none	
USERNAME	admin	
AUTHPASS	•••••	
CONFIRM PASSWORD	•••••	
AUTHTYPE	MD5	
ENCRYPTFLAG	● Yes	○ No
ENCRYPTTYPE	DES	
ENCRYPTPASSWORD	••••	
CONFIRM PASSWORD	••••	

Figure 2.3: Configuring the DT Logins test

7. To know how to configure parameters, refer to **Monitoring the Double-Take Availability Server**.
8. Next, try to signout of the eG administrative interface now, you will be prompted to configure the **Network Interfaces** test. To know more about how to configure this test, refer to the *Monitoring Cisco Router* document.
9. Finally, signout of the eG administrative interface.

Chapter 3: Monitoring the Double-Take Availability Server

eG Enterprise offers a specialized Double-Take Availability monitoring model, which monitors the uptime of, the connections to, and rate at which the Double-Take Availability server performs mirroring and replication operations, and proactively alerts administrators to current and potential deviations from desired performance levels.

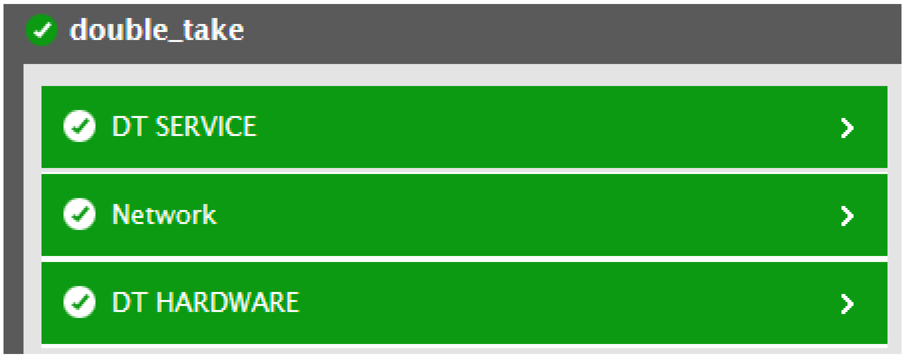

Figure 3.1: Layer model of the Double-Take Availability server

Each layer of Figure 1.2 above is mapped to a wide variety of tests that periodically poll the SNMP MIB of the Double-Take Availability server to capture errors and slowdowns in its functioning. Using the statistics so reported, administrators can infer the following:

- Is the Double-Take Availability server available over the network? If so, how quickly is it responding to requests?
- Are all network interfaces supported by the server operating at normal speeds?
- Is any network interface utilizing bandwidth excessively?
- Is the Double-Take Availability server using too much memory from the reserved memory pool for its operations?
- Was the Double-Take Availability server down recently?
- Has any connection to a target been active for an unusually long time?
- Is a target experiencing any errors in the connections to it?
- Does any connection have too many operations in queue? If so, what type of operations hog the queue - mirroring or replication?
- Have any logins to the Double-Take Availability source and/or target failed? Did any of these login failures occur in the last measurement period?

The sections that follow will discuss each layer of Figure 3.1 in detail.

3.1 The DT HARDWARE Layer

The tests mapped to this layer monitor the uptime and the memory usage of the Double-Take Availability server.

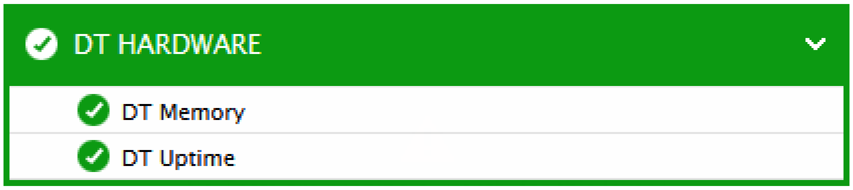

Figure 3.2: Tests mapped to the DT HARDWARE layer

3.1.1 DT Memory Test

When the Double-Take service starts, it reserves a pool of user-addressable memory equal to the Double-Take pagefile size. This reserved pool of memory is the Double-Take pagefile. Although Double-Take has the pool of memory reserved, it only uses what is necessary at any given time. This test reports the amount of memory in the reserved memory pool that the Double-Take Availability server currently uses, and thus enables you to track the memory usage of the Double-Take Availability server.

Target of the test : A Double-Take Availability server

Agent deploying the test : A remote agent

Outputs of the test : One set of results for the Double-Take Availability server being monitored.

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed.
Host	The IP address of the host for which this test is to be configured.
SNMPPort	The port at which the monitored target exposes its SNMP MIB; The default value is 161.
SNMPVersion	By default, the eG agent supports SNMP version 1. Accordingly, the default selection in the SNMPversion list is **v1**. However, if a different SNMP framework is in use in your environment, say SNMP **v2** or **v3**, then select the corresponding option from this list.

Parameter	Description
SNMPCommunity	The SNMP community name that the test uses to communicate with the firewall. This parameter is specific to SNMP **v1** and **v2** only. Therefore, if the SNMPVersion chosen is **v3**, then this parameter will not appear.
UserName	This parameter appears only when **v3** is selected as the SNMPVersion. SNMP version 3 (SNMPv3) is an extensible SNMP Framework which supplements the SNMPv2 Framework, by additionally supporting message security, access control, and remote SNMP configuration capabilities. To extract performance statistics from the MIB using the highly secure SNMP v3 protocol, the eG agent has to be configured with the required access privileges – in other words, the eG agent should connect to the MIB using the credentials of a user with access permissions to be MIB. Therefore, specify the name of such a user against this parameter.
Context	This parameter appears only when v3 is selected as the SNMPVersion. An SNMP context is a collection of management information accessible by an SNMP entity. An item of management information may exist in more than one context and an SNMP entity potentially has access to many contexts. A context is identified by the SNMPEngineID value of the entity hosting the management information (also called a contextEngineID) and a context name that identifies the specific context (also called a contextName). If the Username provided is associated with a context name, then the eG agent will be able to poll the MIB and collect metrics only if it is configured with the context name as well. In such cases therefore, specify the context name of the Username in the Context text box. By default, this parameter is set to *none*.
AuthPass	Specify the password that corresponds to the above-mentioned Username. This parameter once again appears only if the SNMPversion selected is **v3**.
Confirm Password	Confirm the AuthPass by retyping it here.
AuthType	This parameter too appears only if **v3** is selected as the SNMPversion. From the Authtype list box, choose the authentication algorithm using which SNMP v3 converts the specified username and password into a 32-bit format to ensure security of SNMP transactions. You can choose between the following options: • **MD5** – Message Digest Algorithm • **SHA** – Secure Hash Algorithm
EncryptFlag	This flag appears only when **v3** is selected as the SNMPversion. By default, the eG agent does not encrypt SNMP requests. Accordingly, the this flag is set to **No** by default. To ensure that SNMP requests sent by the eG agent are encrypted, select the **Yes** option.
EncryptType	If this EncryptFlag is set to **Yes**, then you will have to mention the encryption type by

Parameter	Description
	selecting an option from the EncryptType list. SNMP v3 supports the following encryption types: • **DES** – Data Encryption Standard • **AES** – Advanced Encryption Standard
EncryptPassword	Specify the encryption password here.
Confirm Password	Confirm the encryption password by retyping it here.
Timeout	Specify the duration (in seconds) within which the SNMP query executed by this test should time out in this text box. The default is 10 seconds.
Data Over TCP	By default, in an IT environment, all data transmission occurs over UDP. Some environments however, may be specifically configured to offload a fraction of the data traffic – for instance, certain types of data traffic or traffic pertaining to specific components – to other protocols like TCP, so as to prevent UDP overloads. In such environments, you can instruct the eG agent to conduct the SNMP data traffic related to the monitored target over TCP (and not UDP). For this, set this flag to **Yes**. By default, this flag is set to **No**.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Memory allocated	Indicates the amount of memory from the reserved memory pool that is currently allocated to the mirroring and/or replication operations of the Double-Take server being monitored.	MB	A very high value could indicate that the the server functions are over-utilizing the memory resources available to them. In order to provide maximum Double-Take system performance, Double-Take servers should be configured with enough RAM to accommodate the maximum Double-Take pagefile (1GB) in addition to the server's other memory needs.

3.1.2 DT Uptime Test

In most production environments, it is essential to monitor the uptime of the Double-Take Availability server, as the DR capability of the applications, data, physical, and virtual servers in such environments relies on the availability of the Double-Take Availability server. By tracking the uptime of this server, administrators can determine what percentage of time Double-Take has been up and the percentage of time it was not. If the data on the source and target servers are not in sync at any given point in time, then, you need to know whether it is because the Double-Take server was down during that time period.

The DT Uptime test included in the eG agent monitors the uptime of the Double-Take Availability server.

Target of the test : A Double-Take Availability server

Agent deploying the test : A remote agent

Outputs of the test : One set of results for the Double-Take Availability server being monitored.

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed.
Host	The IP address of the host for which this test is to be configured.
SNMPPort	The port at which the monitored target exposes its SNMP MIB; The default value is 161.
SNMPVersion	By default, the eG agent supports SNMP version 1. Accordingly, the default selection in the SNMPversion list is **v1**. However, if a different SNMP framework is in use in your environment, say SNMP **v2** or **v3**, then select the corresponding option from this list.
SNMPCommunity	The SNMP community name that the test uses to communicate with the firewall. This parameter is specific to SNMP **v1** and **v2** only. Therefore, if the SNMPVersion chosen is **v3**, then this parameter will not appear.
UserName	This parameter appears only when **v3** is selected as the SNMPVersion. SNMP version 3 (SNMPv3) is an extensible SNMP Framework which supplements the SNMPv2 Framework, by additionally supporting message security, access control, and remote SNMP configuration capabilities. To extract performance statistics from the MIB using the highly secure SNMP v3 protocol, the eG agent has to be configured with the required access privileges – in other words, the eG agent should connect to the MIB

Parameter	Description
	using the credentials of a user with access permissions to be MIB. Therefore, specify the name of such a user against this parameter.
Context	This parameter appears only when v3 is selected as the SNMPVersion. An SNMP context is a collection of management information accessible by an SNMP entity. An item of management information may exist in more than one context and an SNMP entity potentially has access to many contexts. A context is identified by the SNMPEngineID value of the entity hosting the management information (also called a contextEngineID) and a context name that identifies the specific context (also called a contextName). If the Username provided is associated with a context name, then the eG agent will be able to poll the MIB and collect metrics only if it is configured with the context name as well. In such cases therefore, specify the context name of the Username in the Context text box. By default, this parameter is set to *none*.
AuthPass	Specify the password that corresponds to the above-mentioned Username. This parameter once again appears only if the SNMPversion selected is **v3**.
Confirm Password	Confirm the AuthPass by retyping it here.
AuthType	This parameter too appears only if **v3** is selected as the SNMPversion. From the Authtype list box, choose the authentication algorithm using which SNMP v3 converts the specified username and password into a 32-bit format to ensure security of SNMP transactions. You can choose between the following options: • **MD5** – Message Digest Algorithm • **SHA** – Secure Hash Algorithm
EncryptFlag	This flag appears only when **v3** is selected as the SNMPversion. By default, the eG agent does not encrypt SNMP requests. Accordingly, the this flag is set to **No** by default. To ensure that SNMP requests sent by the eG agent are encrypted, select the **Yes** option.
EncryptType	If this EncryptFlag is set to **Yes**, then you will have to mention the encryption type by selecting an option from the EncryptType list. SNMP v3 supports the following encryption types: • **DES** – Data Encryption Standard • **AES** – Advanced Encryption Standard
EncryptPassword	Specify the encryption password here.
Confirm Password	Confirm the encryption password by retyping it here.

Parameter	Description
Timeout	Specify the duration (in seconds) within which the SNMP query executed by this test should time out in this text box. The default is 10 seconds.
Data Over TCP	By default, in an IT environment, all data transmission occurs over UDP. Some environments however, may be specifically configured to offload a fraction of the data traffic – for instance, certain types of data traffic or traffic pertaining to specific components – to other protocols like TCP, so as to prevent UDP overloads. In such environments, you can instruct the eG agent to conduct the SNMP data traffic related to the monitored target over TCP (and not UDP). For this, set this flag to **Yes**. By default, this flag is set to **No**.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Has system been rebooted	Indicates whether the server has been rebooted during the last measurement period or not.	Boolean	If this measure shows 1, it means that the server was rebooted during the last measurement period. By checking the time periods when this metric changes from 0 to 1, an administrator can determine the times when this server was rebooted.
Uptime during the last measure period	Indicates the time period that the system has been up since the last time this test ran.	Secs	If the server has not been rebooted during the last measurement period and the agent has been running continuously, this value will be equal to the measurement period. If the server was rebooted during the last measurement period, this value will be less than the measurement period of the test. For example, if the measurement period is 300 secs, and if the server was rebooted 120 secs back, this metric will report a value of 120 seconds. The accuracy of this metric is dependent on the measurement period – the smaller the measurement period, greater the accuracy.

Measurement	Description	Measurement Unit	Interpretation
Total uptime of the system	Indicates the total time that the server has been up since its last reboot.	Mins	Administrators may wish to be alerted if a server has been running without a reboot for a very long period. Setting a threshold for this metric allows administrators to determine such conditions.

3.2 The Network Layer

Using the tests mapped to this layer, you can verify the network availability of the Double-Take Availability server and measure the speed and bandwidth usage of each of the network interfaces supported by the server.

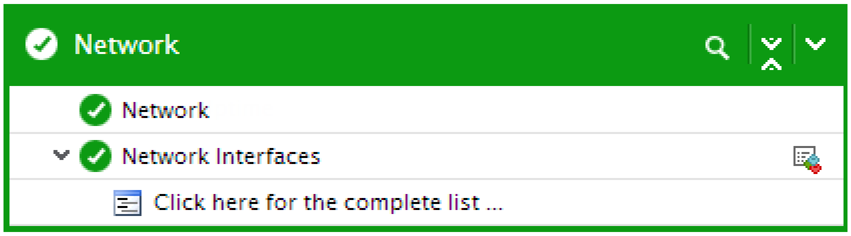

Figure 3.3: The tests mapped to the Network layer

To know how to configure the tests mapped to this layer, refer to in the *Monitoring Cisco Router* document.

3.3 The DT Service Layer

For each target that a replication set connects to, the **DT Connections** test mapped to this layer reports the state of the connection and the level of activity on the connection. In addition, using the **DT Logins** test mapped to this layer, you can accurately point to the unsuccessful login attempts to the Double-Take source and targets.

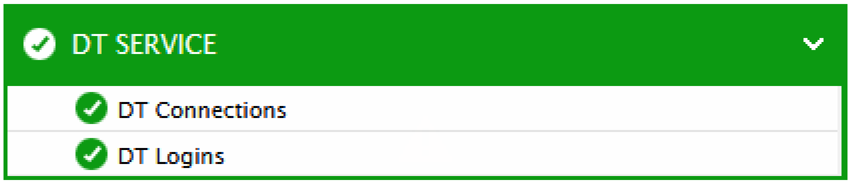

Figure 3.4: The tests mapped to the DT Service layer

3.3.1 DT Connections Test

Protecting specific data consists of two main tasks - creating a replication set (to identify the data to protect) and connecting that replication set to a target. A unique connection ID is associated with each target a replication set connects to. The connection ID provides a reference point for each connection. The connection ID is determined by sequential numbers starting at one (1). Each time a connection is established, the ID counter is incremented. It is reset back to one each time the Double- Take service is restarted. For example, if the Double-Take service was started and the same replication set was connected to five target machines, each connection would have a unique connection ID from 1 to 5.

This test monitors the current state of each Double-Take Availability connection and reports and reports the level of activity on each connection, so that the busiest/overloaded connections are revealed, and the operation (mirroring/replication) that is causing the overload can be identified.

Target of the test : A Double-Take Availability server

Agent deploying the test : A remote agent

Outputs of the test : One set of results for each Double-Take Availability connection being monitored.

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed.
Host	The IP address of the host for which this test is to be configured.
SNMPPort	The port at which the monitored target exposes its SNMP MIB; The default value is 161.
SNMPVersion	By default, the eG agent supports SNMP version 1. Accordingly, the default selection in the SNMPversion list is **v1**. However, if a different SNMP framework is in use in your environment, say SNMP **v2** or **v3**, then select the corresponding option from this list.
SNMPCommunity	The SNMP community name that the test uses to communicate with the firewall. This parameter is specific to SNMP **v1** and **v2** only. Therefore, if the SNMPVersion chosen is **v3**, then this parameter will not appear.
UserName	This parameter appears only when **v3** is selected as the SNMPVersion. SNMP version 3 (SNMPv3) is an extensible SNMP Framework which supplements the SNMPv2 Framework, by additionally supporting message security, access control, and remote

Parameter	Description
	SNMP configuration capabilities. To extract performance statistics from the MIB using the highly secure SNMP v3 protocol, the eG agent has to be configured with the required access privileges – in other words, the eG agent should connect to the MIB using the credentials of a user with access permissions to be MIB. Therefore, specify the name of such a user against this parameter.
Context	This parameter appears only when v3 is selected as the SNMPVersion. An SNMP context is a collection of management information accessible by an SNMP entity. An item of management information may exist in more than one context and an SNMP entity potentially has access to many contexts. A context is identified by the SNMPEngineID value of the entity hosting the management information (also called a contextEngineID) and a context name that identifies the specific context (also called a contextName). If the Username provided is associated with a context name, then the eG agent will be able to poll the MIB and collect metrics only if it is configured with the context name as well. In such cases therefore, specify the context name of the Username in the Context text box. By default, this parameter is set to *none*.
AuthPass	Specify the password that corresponds to the above-mentioned Username. This parameter once again appears only if the SNMPversion selected is **v3**.
Confirm Password	Confirm the AuthPass by retyping it here.
AuthType	This parameter too appears only if **v3** is selected as the SNMPversion. From the Authtype list box, choose the authentication algorithm using which SNMP v3 converts the specified username and password into a 32-bit format to ensure security of SNMP transactions. You can choose between the following options: • **MD5** – Message Digest Algorithm • **SHA** – Secure Hash Algorithm
EncryptFlag	This flag appears only when **v3** is selected as the SNMPversion. By default, the eG agent does not encrypt SNMP requests. Accordingly, the this flag is set to **No** by default. To ensure that SNMP requests sent by the eG agent are encrypted, select the **Yes** option.
EncryptType	If this EncryptFlag is set to **Yes**, then you will have to mention the encryption type by selecting an option from the EncryptType list. SNMP v3 supports the following encryption types: • **DES** – Data Encryption Standard • **AES** – Advanced Encryption Standard

Parameter	Description
EncryptPassword	Specify the encryption password here.
Confirm Password	Confirm the encryption password by retyping it here.
Timeout	Specify the duration (in seconds) within which the SNMP query executed by this test should time out in this text box. The default is 10 seconds.
Data Over TCP	By default, in an IT environment, all data transmission occurs over UDP. Some environments however, may be specifically configured to offload a fraction of the data traffic – for instance, certain types of data traffic or traffic pertaining to specific components – to other protocols like TCP, so as to prevent UDP overloads. In such environments, you can instruct the eG agent to conduct the SNMP data traffic related to the monitored target over TCP (and not UDP). For this, set this flag to **Yes**. By default, this flag is set to **No**.

Measurements made by the test

<table>
<tr><th>Measurement</th><th>Description</th><th>Measurement Unit</th><th>Interpretation</th></tr>
<tr><td>Connection activity</td><td>Indicates the amount of time this connection has been active.</td><td>Minutes</td><td>A very high value for this measure could indicate that the data is taking too long to be transmitted to the target.</td></tr>
<tr><td>Connection state</td><td>Indicates the current state of this connection.</td><td></td><td>The values reported by this measure, their description, and the numeric values that correspond to them have been discussed in the table below:
<table>
<tr><th>Measure Value</th><th>Numeric Value</th><th>Measure Description</th></tr>
<tr><td>conError</td><td>0</td><td>A transmission error has occurred. Possible errors include a broken physical line or a failed target service.</td></tr>
</table>
</td></tr>
</table>

<table>
<tr><th>Measurement</th><th>Description</th><th>Measurement Unit</th><th>Interpretation</th></tr>
<tr><td></td><td></td><td></td><td>
<table>
<tr><th>Measure Value</th><th>Numeric Value</th><th>Measure Description</th></tr>
<tr><td>conActive</td><td>1</td><td>Indicates that the connection is functioning normally and has no scheduling restrictions imposed on it at this time. (There may be restrictions, but it is currently in a state that allows it to transmit.)</td></tr>
<tr><td>conPaused</td><td>2</td><td>Indicates a connection that has been paused. This implies that the network connection exists and is available for data transmission, but the replication and mirror data is being held in a queue and is not being transmitted to the target.</td></tr>
<tr><td>conScheduled</td><td>3</td><td>indicates a connection that is not currently transmitting due to scheduling restrictions (bandwidth limitations,</td></tr>
</table>
</td></tr>
</table>

<table>
<tr><th>Measurement</th><th>Description</th><th>Measurement Unit</th><th>Interpretation</th></tr>
<tr><td></td><td></td><td></td><td><table><tr><th>Measure Value</th><th>Numeric Value</th><th>Measure Description</th></tr><tr><td></td><td></td><td>time frame limitations, and so on).</td></tr><tr><td>conNone</td><td>4</td><td>Indicates that a connection has not been established.</td></tr></table></td></tr>
<tr><td>Ops in Retransmit Queue</td><td>Indicates the number of operations (create, modify, or delete) currently in the retransmit queue on the source.</td><td>Number</td><td></td></tr>
<tr><td>Ops Awaiting Acknowledgments</td><td>Indicates the number of of operations currently waiting in the acknowledgment queue.</td><td>Number</td><td>Each operation that is generated receives an acknowledgment from the target after that operation has been received by the target. This statistic indicates the number of operations that are yet to receive acknowledgment of receipt.</td></tr>
<tr><td>Replication Ops Queued</td><td>Indicates the number of replication operations currently waiting to be executed on the target.</td><td>Number</td><td>Replication is the real-time transmission of file changes to a target. These changes, instead of being replicated to a target, may be queued to disk, if a locked file on the target prevents the changes from being written to it, or if a file on the source changes faster than can be transmitted to the target. Typically, if the system memory allocated to queueing is utilized fully, then new file changes that are to be replicated to a target will be directly queued to disk, while old changes remain in the system memory. Data queued to disk is written to a transaction log.</td></tr>
</table>

Measurement	Description	Measurement Unit	Interpretation
			The value of this measure indicates the number of file changes that are in queue, and are yet to be replicated to the target. A high value of this measure may indicate that too many file changes are awaiting processing or that one/more files on the target have been locked for too long a time.
Mirror Ops Queued	Indicates the number of mirroring operations currently in queue.	Number	Mirroring is the process of transmitting user-specified data from the source to the target, so that an identical copy of data exists on the target. When Double-Take Availability initially performs mirroring, it copies all of the selected data, including file attributes and permissions. Mirroring creates a foundation upon which Double-Take Availability can efficiently update the target server by replicating only file changes. A high value of this measure indicates that many mirroring operations are pending processing, which could hint at a probable processing bottleneck.
Replication Ops Queued Data	Represents the amount of data that was associated with the queued replication operations during the last measurement period.	KB	
Mirror Ops Queued Data	Represents the amount of data that was associated with the queued mirror operations during the last measurement period.	KB	

Measurement	Description	Measurement Unit	Interpretation
Operations Transmitted	Indicates the total number of operations that are currently transmitted to the target.	Number	
Data Sent	Indicates the total number of bytes sent to the target since the last measurement period.	KB	
Operations Received	Indicates the total number of operations (create, modify, or delete) currently received from the target.	Number	
Data Received	Indicates the total number of bytes received from the target during the last measurement period.	KB	
Resent operations	Indicates the number of operations that were resent because they were not acknowledged.	Number	

3.3.2 DT Logins Test

To ensure protection of your data, Double-Take Availability offers multi-level security using native operating system security features. Privileges are granted through membership in user groups defined on each machine running Double-Take Availability. To gain access to a particular Double-Take Availability source or target, the user must provide a valid operating system user name and password and the specified user name must be a member of one of the Double-Take Availability security groups. Once a valid user name and password has been provided and the Double-Take Availability source or target has verified membership in one of the Double-Take Availability security groups, the user is granted appropriate access to the source or target and the corresponding features are enabled in the client.

Using this test, you can promptly detect a failed login attempts to a Double-Take source or target.

Target of the test : A Double-Take Availability server

Agent deploying the test : A remote agent

Outputs of the test : One set of results for each Double-Take Availability server being monitored.

Configurable parameters for the test

Parameter	Description
Test Period	How often should the test be executed.
Host	The IP address of the host for which this test is to be configured.
SNMPPort	The port at which the monitored target exposes its SNMP MIB; The default value is 161.
SNMPVersion	By default, the eG agent supports SNMP version 1. Accordingly, the default selection in the SNMPversion list is **v1**. However, if a different SNMP framework is in use in your environment, say SNMP **v2** or **v3**, then select the corresponding option from this list.
SNMPCommunity	The SNMP community name that the test uses to communicate with the firewall. This parameter is specific to SNMP **v1** and **v2** only. Therefore, if the SNMPVersion chosen is **v3**, then this parameter will not appear.
UserName	This parameter appears only when **v3** is selected as the SNMPVersion. SNMP version 3 (SNMPv3) is an extensible SNMP Framework which supplements the SNMPv2 Framework, by additionally supporting message security, access control, and remote SNMP configuration capabilities. To extract performance statistics from the MIB using the highly secure SNMP v3 protocol, the eG agent has to be configured with the required access privileges – in other words, the eG agent should connect to the MIB using the credentials of a user with access permissions to be MIB. Therefore, specify the name of such a user against this parameter.
Context	This parameter appears only when v3 is selected as the SNMPVersion. An SNMP context is a collection of management information accessible by an SNMP entity. An item of management information may exist in more than one context and an SNMP entity potentially has access to many contexts. A context is identified by the SNMPEngineID value of the entity hosting the management information (also called a contextEngineID) and a context name that identifies the specific context (also called a contextName). If the Username provided is associated with a context name, then the eG agent will be able to poll the MIB and collect metrics only if it is configured with the context name as well. In such cases therefore, specify the context name of the Username in the Context text box. By default, this parameter is set to *none*.

Parameter	Description
AuthPass	Specify the password that corresponds to the above-mentioned Username. This parameter once again appears only if the SNMPversion selected is **v3**.
Confirm Password	Confirm the AuthPass by retyping it here.
AuthType	This parameter too appears only if **v3** is selected as the SNMPversion. From the Authtype list box, choose the authentication algorithm using which SNMP v3 converts the specified username and password into a 32-bit format to ensure security of SNMP transactions. You can choose between the following options: • **MD5** – Message Digest Algorithm • **SHA** – Secure Hash Algorithm
EncryptFlag	This flag appears only when **v3** is selected as the SNMPversion. By default, the eG agent does not encrypt SNMP requests. Accordingly, the this flag is set to **No** by default. To ensure that SNMP requests sent by the eG agent are encrypted, select the **Yes** option.
EncryptType	If this EncryptFlag is set to **Yes**, then you will have to mention the encryption type by selecting an option from the EncryptType list. SNMP v3 supports the following encryption types: • **DES** – Data Encryption Standard • **AES** – Advanced Encryption Standard
EncryptPassword	Specify the encryption password here.
Confirm Password	Confirm the encryption password by retyping it here.
Timeout	Specify the duration (in seconds) within which the SNMP query executed by this test should time out in this text box. The default is 10 seconds.
Data Over TCP	By default, in an IT environment, all data transmission occurs over UDP. Some environments however, may be specifically configured to offload a fraction of the data traffic – for instance, certain types of data traffic or traffic pertaining to specific components – to other protocols like TCP, so as to prevent UDP overloads. In such environments, you can instruct the eG agent to conduct the SNMP data traffic related to the monitored target over TCP (and not UDP). For this, set this flag to **Yes**. By default, this flag is set to **No**.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Total Successful Logins	Indicates the total number of successful logins to the server.	Number	
Total Failed Logins	Indicates the total number of failed logins to the server.	Number	Ideally, the value of this measure should be 0.
Current Successful Logins	Indicates the number of login attempts that were successful during the last measurement period.	Number	
Current Failed Logins	Indicates the number of login attempts that failed during the last measurement period.	Number	Ideally, the value of this measure should be 0.

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.

This document may not be reproduced by any means nor modified, decompiled, disassembled, published or distributed, in whole or in part, or translated to any electronic medium or other means without the prior written consent of eG Innovations. eG Innovations makes no warranty of any kind with regard to the software and documentation, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. The information contained in this document is subject to change without notice.

All right, title, and interest in and to the software and documentation are and shall remain the exclusive property of eG Innovations. All trademarks, marked and not marked, are the property of their respective owners. Specifications subject to change without notice.