Microsoft Dynamics NAV

eG Enterprise v6

Table of Contents

Table of Figures

1

Monitoring Microsoft Dynamics NAV

Microsoft Dynamics NAV is an enterprise resource planning (ERP) software product from Microsoft. It automates and streamlines business processes for small and mid-sized organizations and local subsidiaries of large international Groups. The product is part of the Microsoft Dynamics family, and is intended to assist with finance, manufacturing, customer relationship management, supply chains, analytics and electronic commerce.

The Microsoft Dynamics NAV software is composed of three major components (three-tier architecture):

- **Data Tier**: The Database Server, a database that stores the Microsoft Dynamics NAV data
- **Middle or Server Tier**: The Application Server, which is a service that manages all business logic and communication of Microsoft Dynamics NAV's operation
- **Client Tier**: The Client(s) is the actual user interface into Microsoft Dynamics NAV. Through the Microsoft Dynamics NAV Portal Framework, NAV 2013 supports the following clients:

  - The Windows Client

  - The Web Client

  - The Web Service Client

  - The SharePoint client

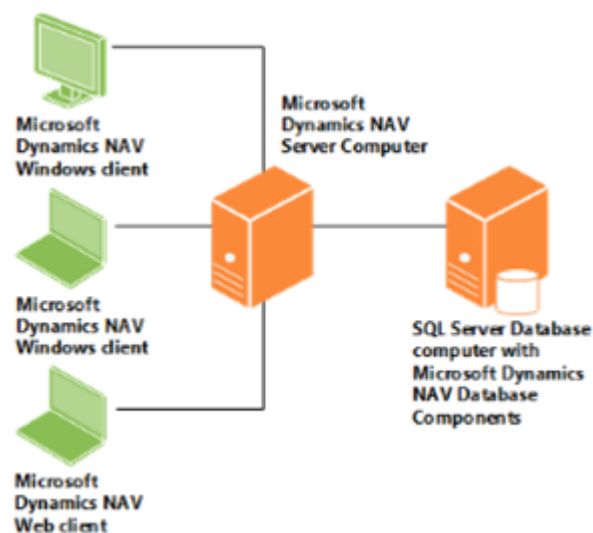 It additionally supports a NAS services client for programmatic access.

Figure 1.1: Architecture of Microsoft Dynamics NAV

From Figure 1.1, it is evident that the **Microsoft Dynamics NAV Server** is critical to the uninterrupted operations of the NAV solution. If this component slows down or is unreachable, then clients will not be able to access critical business data stored in the NAV database whenever they need it! This in turn may cause crucial business processes to stall, resulting in significant loss of revenues and opportunities! If this is to be avoided, administrators must make sure that the Microsoft Dynamics NAV server is always available, is sized right, and is able to service the requests it receives promptly. For this, administrators must continuously monitor the availability and operations of the NAV server, quickly detect bottlenecks, accurately isolate their source, and rapidly resolve them.

For this purpose, eG Enterprise provides a specialized *Microsoft Dynamics NAV* monitoring model. This model monitors the **Microsoft Dynamics NAV Server** computer top-down and reports on its availability and overall health.
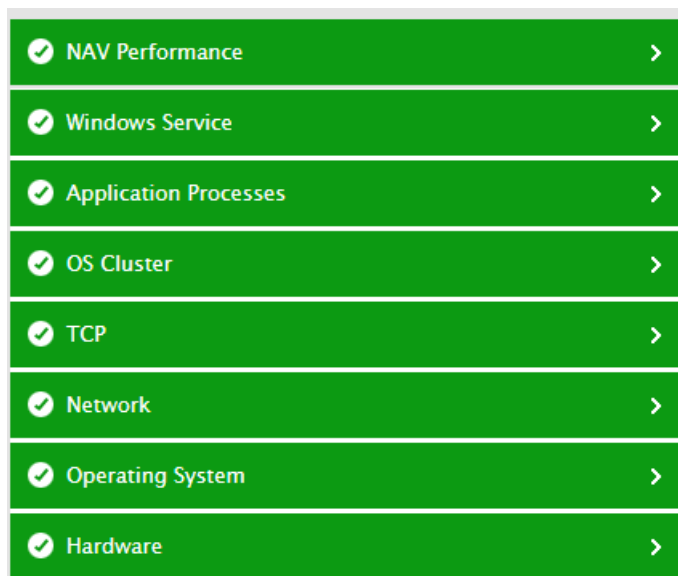
Figure 1.2: Layer model of the Microsoft NAV server

The *Microsoft Dynamics NAV* monitoring model can monitor only **Microsoft Dynamics NAV 2013**.

Each layer in this model is mapped to tests that keep tabs on the health of the server host, the quality of network connections to the host, and the efficiency of the internal operations of the NAV server.

For collecting metrics on the internal health of the NAV server in real-time, eG's agent-based/agentless monitor relies on the Windows Communication Foundation (WCF); this includes a number of performance counters that track requests to the NAV server and measure how quickly the server processes these requests. To enable the eG agent to use WCF performance counters, follow the steps below:

1. Edit the **Microsoft.Dynamics.Nav.Server.exe.config** file in the NAV Service folder (default is *C:\Program Files\Microsoft Dynamics NAV\60\Service*).

2. Add the section below, marked in **Bold**:

```xml
<?xml version="1.0" encoding="utf-8" ?>
<configuration>
  <appSettings file="CustomSettings.config" />
  <system.diagnostics>
    <assert assertuienabled="false" />
  </system.diagnostics>

  <system.serviceModel>
      <diagnostics performanceCounters="All" />
  </system.serviceModel>

</configuration>
```

3.   Restart the NAV Server.

With the help of the metrics so collected, administrators can find quick and accurate answers for the following performance queries:

- Is the current session load on any Nav server instance high?

- Is any NAV server instance taking too long to process server operations? If so, which instance is it?

- Is any server instance re-executing queries on the NAV database?

- Which NAV server instance is taking too long to write to the NAV database?

- Are all NAV server instances using the preferred connection cache optimally?

- Are hits to the Calculated Fields cache very low?

- Have the commands cache, primary key cache, and the result set cache been utilized well?

- Are too many calls to the NAV service pending processing?

- Are calls failing frequently?

- At what rate do calls returns faults?

- Are calls to the service taking too long to be processed?

- Have any messages been dropped by the queued transport?

- How many messages were rejected by the queued transport?

- Were any messages received from unauthorized users?

- What is the average transaction load on the NAV service?

- Are these transactions getting aborted frequently?

The sections that follow will discuss the topmost layer of Figure 1.1 alone, as the rest of the layers have already been discussed in the *Monitoring Unix and Windows Servers* document.

1.1 The NAV Performance Layer

This layer reports metrics on the performance of the NAV server and service.

1.1.1 NAV Server Performance Test

If a Microsoft Dynamics NAV server instance is not able to process client requests quickly, administrators must be able to rapidly detect the slowdown and precisely pinpoint its root-cause, well before users notice and complain. One of the reasons for this slowness could be the poor processing power of the NAV server itself. Another key reason for this slowdown, but one which is often ignored, is the poor performance of the Microsoft SQL server that hosts the NAV data. If this database server and its caches are not configured right or if queries to this server are constructed inefficiently, then an NAV server instance will not be able to update or retrieve data from the database swiftly; this in turn will cripple the server's processing ability and will result in a slowdown in NAV server operations.

In such cases, if the administrator monitors only the NAV server instance, and ignores server-database interactions, he/she may not be able identify the exact reason for the slowdown. As an outcome, troubleshooting cycles will grow longer and problem resolution will be delayed. This is where the **NAV Server Performance** test helps! This test tracks client sessions to every NAV server instance, continuously measures the time taken for server operations, and thus enables administrators to spot any slowness in the operations and the instance that has been impacted by this. By additionally keeping an eye on the server-database interactions, the test reveals whether/not the SQL database server is contributing to the slowness, and if so, in which way.

Purpose	Tracks client sessions to every NAV server instance, continuously measures the time taken for server operations, and thus enables administrators to spot any slowness in the operations and the instance that has been impacted by this. By additionally keeping an eye on the server-database interactions, the test reveals whether/not the SQL database server is contributing to the slowness, and if so, in which way
Target of the test	A Microsoft Dynamics NAV Server
Agent deploying the test	An internal/remote agent
Configurable parameters for the test	1. **TESTPERIOD** – How often should the test be executed 2. **HOST** – The IP address of the NAV server 3. **PORT** – The port at which the NAV server listens.

Outputs of the test	One set of results for every NAV server instance being monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Active sessions: Indicates the number of active sessions on this Microsoft Dynamics NAV Server instance.	Number	An active session is a connection to the Microsoft Dynamics NAV Server instance from a Microsoft Dynamics NAV client, such as the Microsoft Dynamics NAV Windows client or Microsoft Dynamics NAV Web client, NAS, or Web services. This measure is a good indicator of the current session load on the NAV server.
	Server operations rate: Indicates the number of operations that have started on this server instance per second.	Operations/sec	An operation is a call to the Microsoft Dynamics NAV Server instance from a Microsoft Dynamics NAV client to run Microsoft Dynamics NAV objects. OData and SOAP requests are not included in the value of this measure.
	Average server operations rate: Indicates the average duration of server operations on this instance.	Secs	A high value or a consistent increase in the value of this measure is a definite sign of a processing bottleneck on the server.
	Mounted tenants: Indicates the number of tenants that are mounted on this instance.	Number	Microsoft Dynamics NAV supports deployments where several different companies access a centrally maintained Microsoft Dynamics NAV application. By using this *multitenancy* support, you can add new customers to your solution easily, and you can roll out updates quickly with limited downtime for your customers. A tenant is an entity that uses your solution and stores data in a business database. This is often either a business or a group of legal entities whose data can be stored in one database. In practical terms, a tenant is a database that stores business data for one or more Microsoft Dynamics NAV companies. Each tenant is connected to a Microsoft Dynamics NAV Server instance, but the Microsoft Dynamics NAV Server instance can support multiple tenants.

	Open connections to NAV server: Indicates the current number of open connections from this instance to the Microsoft Dynamics NAV databases on SQL Servers.	Number	
	Percentage of queries re-executed: Indicates the percentage of queries that are re-executed by this instance when fetching the query result.	Percent	A high value for this measure is a cause for concern, as it indicates that many queries were re-executed. This can in turn can delay request processing by the NAV server.
	Hearbeat rate: Indicates the time that it takes for this instance to complete a single write to a system table.	Secs	Every 30 seconds, the Microsoft Dynamics NAV Server instance writes a record to indicate that the instance is "alive". If the NAV server takes too much time to perform this write operation, it indicates a highly latent network connection between the server and the SQL database. Such a slow network connection can also significantly delay NAV server operations.
	Percentage of hits in preferred connection cache: Indicates the percentage of hits in the preferred connection cache, compared to the total number of requests made by this instance to that cache.	Percent	The preferred connection cache contains connections from the SQL connection pool that was last used by a Microsoft Dynamics NAV user. This measure indicates the percentage of requests for SQL connections that were serviced by the connections already present in the preferred connection cache. A high value is desired for this measure as it means that connection requests were serviced by the cache itself, and not by the SQL connection pool. This indicates that the NAV server was able to connect to the SQL database server quickly.

| | **Percentage of hits in calculated fields cache:**

Indicates the percentage of requests made by this instance that were serviced by the calculated fields cache. | Percent | The calculated fields cache contains the results of the CALCFIELDS Function (Record) calls. The CALCFIELDS Function calculates the FlowFields in a record. FlowFields are virtual fields. The values in these fields are not saved in the table. This means that you must use either the CALCFIELDS function or the SETAUTOCALCFIELDS Function (Record) to update them. For example, if you retrieve a record using the FIND Function (Record) and NEXT Function (Record), the FlowFields in those records are set to zero (0). Then, when you call CALCFIELDS, their values are updated.

A high value for this measure implies that a large number of requests for FlowFields were serviced by the calculated fields cache itself, thus minimizing direct database accesses and related overheads. A high value is hence desired for this measure.

A low value on the other hand means that the database had to be accessed to service a majority of requests for FlowFields. This can happen even the calculated fields cache does not have enough space to hold on to the results of the CALCFIELDS Function (Record) calls. You may have to size the cache right to avoid this. |
| | **Percentage of hits in command cache:**

Indicates the percentage of requests made by this instance that were serviced by the SQL command cache. | Percent | The command cache contains the results of all SQL commands.

If the results of a SQL command are available in the commands cache, then a query need not retrieve the same by directly accessing the database. This significantly improves database performance, as it greatly reduces expensive database accesses. This is why, the value of this measure should ideally be high.

A low value on the other hand implies that many requests did not find their results in the command cache, and were hence forced to access the database for the same. This can happen if the command cache is not sized right, and is hence not able to store many SQL command results. If this is so, then you may want to consider fine-tuning the cache size. |

	Percentage of hits in primary key cache: Indicates the percentage of requests made by this instance that were serviced by the primary key cache.	Percent	The primary key cache contains the results of requests to get a record by using its primary key. A high value for this measure indicates that many queries that use the primary key for retrieving records are being serviced by the primary key cache itself, and are not being sent to the database for processing. As this reduces expensive database accesses, it is good practice to keep the value of this measure high. A low value for this measure on the other hand implies that a majority of such queries are being processed in the SQL database only, as the primary key cache does not have the desired results. This could be because enough memory has not been assigned to the cache. You may want to consider resizing the cache for better results.
	Percentage of hits in result set cache: Indicates the percentage of requests made by this instance that were serviced by the result set cache.	Percent	The result set cache contains result sets that are returned from SQL Server. A well-sized result set cache can significantly improve database performance by reducing direct database accesses for servicing requests. In this case, the value of this measure will be high. A low value on the other hand could imply that the cache is badly sized, rendering it unable to handle requests. This is bound to have an adverse impact on database health, as it means that requests are accessing the database often for result sets.
	Number of rows in all temporary tables: Indicates the number of rows in all temporary tables for this instance.	Number	

1.1.2   NAV Service Performance Test

To fully understand how well a Microsoft Dynamics NAV server instance is performing and to proactively detect potential glitches in its operations, an administrator has to:

- Closely track the calls made to and trasactions executed by the NAV service of instance;

- Determine how well the service is handling these calls and transactions, and;

- Accurately isolate processing bottlenecks/failures/faults before users do.

The **NAV Service Performance** test helps with this. This test observes the call and transaction load on every instance of the NAV service and checks whether/not that service instance is able to handle this load. In the process, pending calls, failed calls, faulted calls, and aborted transactions are brought to the attention of the administrator, thus alerting him/her to potential problems in NAV operations.

Purpose	Observes the call and transaction load on every instance of the NAV service and checks whether/not that service instance is able to handle this load. In the process, pending calls, failed calls, faulted calls, and aborted transactions are brought to the attention of the administrator, thus alerting him/her to potential problems in NAV operations		
Target of the test	A Microsoft Dynamics NAV Server		
Agent deploying the test	An internal/remote agent		
Configurable parameters for the test	1. **TESTPERIOD** – How often should the test be executed 2. **HOST** – The IP address of the NAV server 3. **PORT** – The port at which the NAV server listens.		
Outputs of the test	One set of results for every instance of the NAV service being monitored		
Measurements made by the test	**Measurement**	**Measurement Unit**	**Interpretation**
	Calls: Indicates the number of calls to this service.	Number	This is a good indicator of the workload of the NAV service.
	Calls rate: Indicates the rate at which calls are made to this service.	Calls/sec	
	Calls outstanding: Indicates the number of in-progress calls to this service.	Number	A consistent increase in the value of this measure indicates that the service is unable to process calls as quickly as it receives them. If the situation persists, the NAV server may choke due to overload. You may want to allocate more processing power to the server to ensure speedy servicing of calls.
	Calls failed rate: Indicates the rate of calls that have unhandled exceptions, and are received by this service.	Failed/sec	A low value is desired for this measure.

	Calls faulted rate: Indicates the rate at which calls have returned faults to this service.	Faulted/sec	A low value is desired for this measure.
	Call duration: Indicates the average duration of calls to this service.	Secs	A steady/sporadic spike in the value of this measure is a matter of concern, as it indicates a processing bottleneck.
	Instance creation rate: Indicates the total number of service instances created in a second.	Created/sec	
	Queued messages dropped rate: Indicates the number of messages that were dropped by the queued transport at this service per second.	Dropped/sec	Ideally, the value of this measure should be low.
	Queued messages rejected rate: Indicates the number of messages that are rejected by the queued transport at this service every second.	Rejects/sec	Ideally, the value of this measure should be low.
	Unauthorized security calls rate: Indicates the rate at which messages are received from a valid user and are protected properly, but the user is not authorized to do specific tasks.	Calls/sec	A low value is desired for this measure. If the frequency of these calls increases, you may want to investigate the reasons.
	Authentication failures rate: Indicates the rate at which messages are rejected due to a security problem other than unauthorized calls.	Failures/sec	A low value is desired for this measure. If the frequency of failures increases, you may want to investigate the reasons.

	Transactions flowed rate: Indicates the rate at which transactions flowed to operations in this service.	Trans/sec	
	Transactional operations aborted rate: Indicates the rate at which transactional operations were aborted in this service	Operations/sec	Ideally, the value of this measure should be very low.
	Transactional operations committed rate: Indicates the rate at which transactional operations are committed in this service.	Operations/sec	

2

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 the **Microsoft Dynamics NAV** 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.