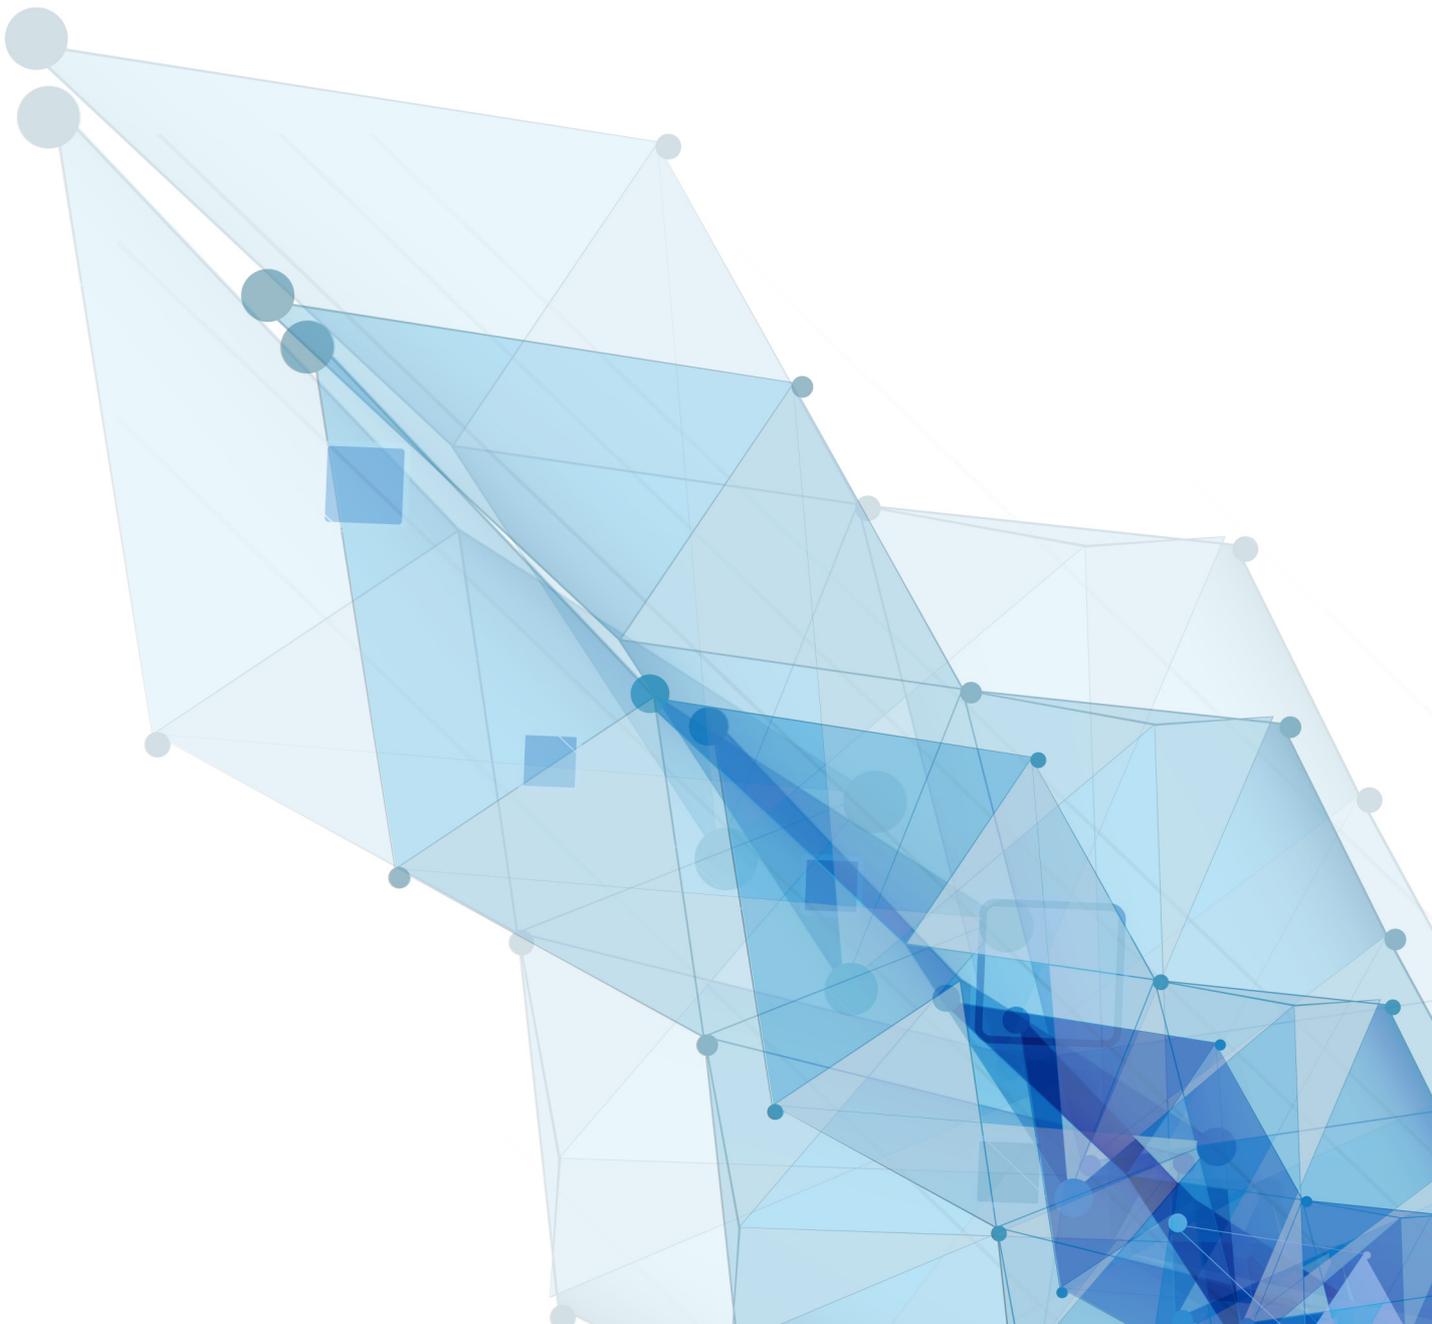
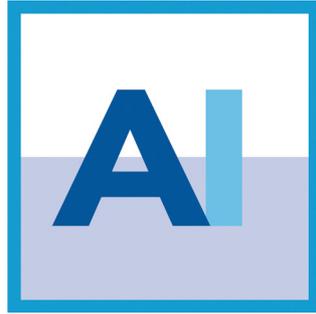


Application Insights™

User Guide





USER GUIDE

Contact

panagenda Austria

(Headquarters)

panagenda GmbH
Schreyvogelgasse 3/10
AT 1010 Vienna (Austria)
Phone: +43 1 89 012 89
Fax: +43 1 89 012 89 – 15

panagenda Germany

panagenda GmbH
Lahnstrasse 17
DE 64646 Heppenheim
(Germany)
Phone: +49 6252 67939 – 00
Fax: +49 6252 67939 – 16

panagenda USA

panagenda Inc.
60 State Street
Suite 700
Boston, MA 02109 (USA)
Phone: +1 (617) 855 5961
Fax: +1 (617) 488 2292

E-Mail Sales: sales@panagenda
E-Mail Support: support@panagenda
Web: www.panagenda

Table of Contents

1	General Information	5
1.1	Welcome to panagenda ApplicationInsights!	5
1.2	What is panagenda ApplicationInsights?	5
1.3	Key use cases for panagenda ApplicationInsights	6
1.3.1	Optimization	6
1.3.2	Modernization	6
1.3.3	Transformation	7
1.4	Licensing & Entitlement	8
1.5	Client system requirements	9
2	Data Collection	10
2.1	Data Collection Scope	10
2.2	Instances versus Databases	10
2.3	Definitions	11
2.3.1	Determination of Database Types & Focus Databases	11
2.3.2	User access days versus sessions	12
2.3.3	Complexity, Design impact, Design similarities and Design Insights	13
2.4	Data Collection and Time Constraints (Servers Being Polled/ Period Collected)	13
2.4.1	Why don't I see all info?/Why is part of the graph greyed out?	14
3	Getting Started with ApplicationInsights	15
3.1	Login	15
3.2	User Roles	15
3.3	System Settings	16
3.4	Using panagenda ApplicationInsights: Interface Tips	16
3.4.1	Selecting/Deselecting subsets	17
3.4.2	Maximize screen usage	17
3.4.3	Help and additional information	18
4	AI Components	19
4.1	Overview	19
4.1.1	Status Dashboard	19
4.1.2	Environment Overview	20
4.1.3	Inventory	21

4.2 Catalog	21
4.2.1 Working with the Catalog List	22
4.2.2 Instance Details - General	23
4.2.3 Instance Details - Tab Application	24
4.2.4 Instance Details – Tab Other Instances (x)	24
4.2.5 Instance Details – Tab Usage	25
4.2.6 Instance Details – Tab Design Complexity	25
4.2.7 Instance Details – Tab Design Similarity	25
4.2.8 Instance Details – Tab Design Insights	26
4.3 Usage	28
4.3.1 Sessions versus Complexity	28
4.3.2 Web versus Client Usage	29
4.3.3 Notes Client Usage	30
4.3.4 Web Browser Usage	31
4.4 Design	31
4.4.1 Complexity	31
4.4.2 Similarity	32
4.4.3 Insights	33
5 Design Complexity	35
5.1 Collected Design Data	35
5.2 Design Complexity and Impact Criteria	35
5.3 Weighted analysis	36
5.4 Design Complexity Ruleset	37
6 Support & Help	38
DISCLAIMER	39
Limitation of liability for external links	39

1 General Information

1.1 Welcome to panagenda ApplicationInsights!



This document will help you understand what panagenda ApplicationInsights is, how it works and how you can use it within your organization.

1.2 What is panagenda ApplicationInsights?

panagenda ApplicationInsights (also referred to as “AI” in this document) is delivered as a virtual appliance that provides continuous and detailed information on IBM Notes / Domino application usage and design in the form of easily understandable graphs and visualizations. It offers reports on design complexity of applications, gives insights into the application landscape for optimization, migration, modernization and transformation purposes. panagenda ApplicationInsights offers administrators, developers and management the means to plan, manage execute and monitor a Domino application modernization strategy.

In order to do so, AI provides:

- Environment and inventory overviews
- Insights into usage from IBM Notes Rich clients and browsers
- Analytics on design and source code complexity
- Detailed „Design Insights” based on source code search patterns
- Design similarity and identification of „Template Candidates”
- Continuous collection of data (overview, usage and design changes)

1.3 Key use cases for panagenda ApplicationInsights

There are several use cases in which you will find that panagenda ApplicationInsights can help. The three key areas are:

- Optimization
- Modernization
- Transformation

1.3.1 Optimization

Reduce costs by optimizing the number of applications and servers and by archiving or sunsetting applications, or instances of applications, no longer needed.

Removing unused applications and database instances reduces load, improves maintainability and improves the cost-benefit ratio.

AI can help you, for instance, to identify archiving, sunsetting and transformation candidates in the following ways:

- If an application is mostly used read-only: you could for instance choose to archive, sunset or perform a content migration (to IBM Connections, for example)
- If an application is not used at all (after, say, 6-12 months): You could choose to close down access to the database and if no-one complains after e.g., 3 months, archive it.

1.3.2 Modernization

By providing insights on usage and design, AI allows you to make founded choices about enabling select applications for mobile and/or browser use.

For instance to refurbish select applications used by a large amount of users or a high-profile group of users and departments (e.g. board of management, sales department, profit centers, etc).

To help in this modernization process, AI helps to answer questions like:

- Which applications are worth investing in? Usually those showing high activity or are used by key departments.
- If invested in, which applications pay back fast? Applications with exceptionally complex design may be expensive to modernize, for example, so focus on lesser complexity.
- How can I optimize my modernization investment? Applications that share a lot of design similarity with others, pay back faster, as – having modernized one, could implicitly also modernize all mostly similar applications.
- Which applications depend on backend interfaces on the server side? (e.g. integration with SAP, dependency on fax or file system, etc. are not compatible with a hosted/cloud environment)
- Which applications use UI elements only supported in the Notes Client? (a variety of code and design elements are not supported in ICAA/Browser Plugin or web)

1.3.3 Transformation

Making your environment a more maintainable and more importantly, shapeable Domino environment facilitates adoption of Verse, Connections, Watson Workspace et al.

Using the above mentioned options you can achieve a better form of control over the platform and allow for integration and expansion with additional systems and services.

AI has a lot to offer when determining the forward path for each application:

- Investigate the best approach to optimize and modernize your application landscape
- Identify unused applications (deletion candidates)
- Identify read-only/read-mostly applications (content migration candidates)
- Identify archiving and sun-setting candidates
- Identify migration candidates for IBM Connections or other content platforms
- Identify most valuable applications from both a usage and complexity standpoint
- Identify potential roadblocks in applications (e.g. applications with a lot of java code)

In situations where the focus is on email Cloud migration it allows you to:

- Analyze databases with migration relevant mail dependencies
- Get insight into deviations from the standard template and identify possible roadblocks or co-existence topics

1.4 Licensing & Entitlement

panagenda ApplicationInsights is licensed on a per database instance basis.

For existing IBM customers with valid maintenance for IBM Domino (including IBM Domino Express offerings), a free entitlement license is available to unlock the 50 most used and 50 most complex application database instances together with extensive insights into security, deployment integrity and load patterns for the entire Domino environment.

With the entitlement comes:

- Entitlement to IBM global support
- 12 months usage of the panagenda ApplicationInsights tool from date of release
- Analysis of your entire Domino environment and details on:
 - Environment, Inventory and Usage overviews across all database instances
 - Usage details for 50 most used database instances
 - Design complexity of 50 most complex database instances
 - Design insights for 50 database instances across the most used and most complex database instances

Please check IBM Passport Advantage (see <https://ibm.co/1OxxqRC>) for information and conditions. Additional packages can be bought per 500 database instances. Please contact your IBM reseller or representative for more information.

1.5 Client system requirements

panagenda ApplicationInsights server installation requirements are listed in the technical Setup guide which can be found here: https://www.panagenda.com/download/ApplicationInsights/ApplicationInsights_SetupGuide_EN.pdf

The panagenda ApplicationInsights web interface is compatible with many HTML5 capable devices and browsers. Recommended browsers are Chrome (64 bit) and Firefox (64 bit).

Browser security and Network Access:

No special web browser security settings are required to run the panagenda ApplicationInsights web interface.

To access the panagenda ApplicationInsights web interface, you need to have access to the panagenda ApplicationInsights appliance via TCP/IP, port 80 (HTTP) and port 443 (HTTPS).

2 Data Collection

2.1 Data Collection Scope

ApplicationInsights collects data on an ongoing basis about all the databases within your Domino application environment. Data is retrieved from various sources including:

- Domino Console
- Notes.ini
- Log.nsf
- Domlog.nsf
- Catalog.nsf
- Domino directories including directory assistance
- Design data from the analyzed databases themselves

Information is collected for all database instances available on the servers specified during configuration/setup. Which servers are included in the collection process is specified in the Discovery settings of panagenda ApplicationInsights.

The data is stored on the panagenda ApplicationInsights Appliance indefinitely¹ to allow for trends and overviews over time.

2.2 Instances versus Databases

ApplicationInsights looks at individual instances when collecting data. An instance is defined as a nsf/ntf container living on a server having a replica ID. A Domino application or 'database' is determined by a unique replica id and can be represented by more than one instance (replicas).

Why are we collecting info about different instances instead of just one?

1. Depending on the availability of sufficient diskspace.

By looking at the instances we can give you more granular and detailed information about the application. Like on which servers it exists, which instances are being used by which users and if there are differences in design between replicas (for instance when replication is broken).

It also allows us to indicate differences in design between application instances and problems with replication, the design task, design inheritance settings and duplicate templates that could cause differences in design.

Database / application:

If in this document the term 'database' or in some cases 'application'² is used it represents a unique database identified as a file with extension .NSF or .NTF and all its analyzed³ instances (replicas) on other servers. Also, if in this document the term 'application' is used in reference to analyzed databases, it refers to databases of both type Mail-In databases and Applications.

Instance:

If the term instance is used a single replica of that database is meant. Note that more than one instance with the same replica ID may exist on any given server.

2.3 Definitions

2.3.1 Determination of Database Types & Focus Databases

The appliance automatically categorizes a found NSF/NTF instance into one of the following types of categories:

1. User Mail database (identified by a person document pointing to the instance)
2. Mail-In database (identified by a non-system mail-in document ...)
3. System database (identified via well-known file or template name)
4. mail.box database (identified from Domino server configuration settings)
5. Directory (identified from Domino server configuration/used template)
6. Orphaned User Mail (user mail files no longer associated with a person document)
7. Application databases (all remaining not categorized into cat 1-6)

2. The term 'application' is also used to indicate a category type of databases as defined in the topic "Definitions" on page 11.
3. As panagenda ApplicationsInsights only has insights into the instances on servers it is analyzing it is important to ensure all servers are included in your discovery set. For details about this please refer to the technical documentation.

If a NSF/NTF instance is categorized once on one server, all existing instances with the same replica ID get categorized into the same category.

Based on the identified types, the DBs are divided into two groups: „**Focus Databases**“ (consisting of Applications and Mail-in Dbs) and „**Other Databases**“ (consisting of the remaining database types).

The majority of detailed usage and design analysis is performed on Focus DBs, but there is a variety of charts and reports with information about the other types.

Additional information outside the scope of ApplicationNights regarding these other DB types and infrastructure topics is available in panagenda iDNA. Please check our website or contact us for additional information.

2.3.2 User access days versus sessions

Each time a user accesses an application a session is set up. During the day users can have multiple sessions, as sessions get closed, users log out/login again or because they use multiple windows or clients. So, one user accessing an application over the day can result in multiple sessions. Sessions can also occur when servers and processes access database instances.

A User Access Day is identified as a unique user accessing an application on any given day within the analyzed period.

So, there could be 20 sessions for User A accessing application Z in a period of 7 days but if that user only accessed the application on day 1 and 5 of that period with 14 sessions occurring on the first day and 6 on the fifth day the User Access Days number will still only be 2. In several graphs this distinction is used to eliminate the risk of skewed numbers and to allow you to make a more balanced decision when it comes to usage numbers.

2.3.3 Complexity, Design impact, Design similarities and Design Insights

Design Complexity is based on the magnitude of design elements and the amount of source code in these elements. The calculation is refined by several considerations like code language type, usage of Reader/Author fields, custom XPages control, etc.

Design Similarity analysis identifies how much alike a database is to all other analyzed databases. Databases with high similarity are grouped into design clusters, which allows for quick identification of replicas that might deviate from a common design. It also provides means to identify 'template candidates' for databases that have no template specified.

Design Insights are combinations of certain patterns occurring in the source code. These combinations are called "Findings" and can be anything from identifying platform dependencies when local DLL's are referenced to understanding if a piece of code interacts with other databases.

Design Impact is a categorization of seven levels from "Insignificant" to "Exceptional" that indicate the impact redesigning or migrating the application would have. The Design impact is calculated from three main sources: Complexity, Insights and Similarity.

2.4 Data Collection and Time Constraints (Servers Being Polled/ Period Collected)

AI collects data continuously from the environment and stores it in the virtual appliance's database system. This ensures that the data collected stays available even after server logs have been cleaned up and allows handling and analysis of much larger amounts of data than what would be possible on the Domino platform.

Right after the initial setup of AI, data collection begins and user activity data is collected from that time on. The virtual appliance goes through three steps until the first data is available four days after the initial setup: Usage Collection, Design Complexity analysis and Design Insights analysis. Once at least a part of the Focus DBs have gone through all stages, the AI interface becomes available. Until then, a status overview screen will provide information about the current collection stage of the appliance.

In the AI web interface, a small progress and status indicator on the top right of the content area provides information about the current status.

To avoid spiky load on collected servers, the AI appliance spreads out collection over the day. Usage data has to be collected on a continuous basis though, so constant activity of the Notes ID used by AI is to be expected.

2.4.1 Why don't I see all info?/Why is part of the graph greyed out?

Much of what is displayed in AIs reports and charts is determined by the installed license key. The IBM Entitlement license key will effectively provide various information on somewhere between 50 and 100 database instances (see chapter Licensing and Entitlement for details).

Certain graphs and reports will show data for all database instance regardless of the license key, but certain details might be omitted. Other elements will not show unlicensed content at all.

Please refer to "Licensing & Entitlement" on page 8 to see how your organization can obtain additional licenses.

3 Getting Started with ApplicationInsights

The following section gives info on the most common operations within panagenda ApplicationInsights.

3.1 Login

Access to panagenda ApplicationInsights is on a per user basis. It requires a username and password to login.

Keep in mind that both username and password are case sensitive!

User accounts can be created by the administrator. Please refer to the technical documentation for information on how to reset passwords and create user accounts.

3.2 User Roles

panagenda ApplicationInsights has two types of users:

- **Viewer:** Can access all analysis and metrics data and interact with the visualizations but is restricted from changing panagenda ApplicationInsights system settings.
- **Administrator:** Can access all analysis and metrics data and interact with the visualizations and is allowed to manage panagenda ApplicationInsights system settings. Including creating accounts and resetting user passwords.

3.3 System Settings

System settings are only available to users with the “Administrator” role and consists of several items accessible through the menu in the top right corner of panagenda ApplicationInsights (‘cog wheel’ symbol).

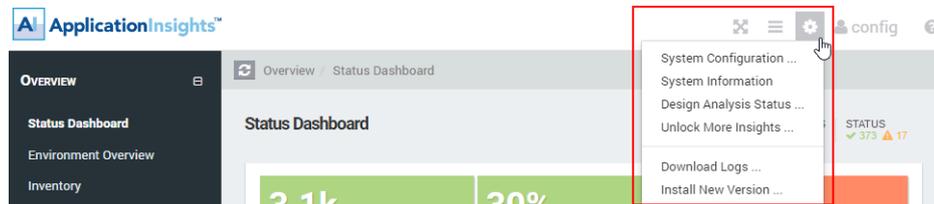


Figure 1: System Settings

- **System configuration:** Links back to the system configuration menu. Please refer to the technical documentation for further details.
- **System Information:** Provides information about the version of AI, the server (Virtual Appliance) it is running on, the browser currently being used to access panagenda ApplicationInsights and the cache. This information can also be requested of you by support in case of problems.
- **Design Collection Status:** Detailed information on design collection status per database instance
- **Unlocking More Insights:** How to obtain additional AI Value Packs
- **Download logs:** Link to the administration section to download logs
- **Install New Version:** Links to the administration section where upgrades can be installed.

3.4 Using panagenda ApplicationInsights: Interface Tips

The various graphs and visualizations within panagenda ApplicationInsights allow you to easily focus in on certain interest areas and data subset. For this several features exist to help you interact with the graphs.

3.4.1 Selecting/Deselecting subsets

Pie charts, bar charts, line graphs and bubble charts that show you multiple categories often allow you to select/deselect certain categories to focus in on subset. To do so, simply click on the category in the charts legend.

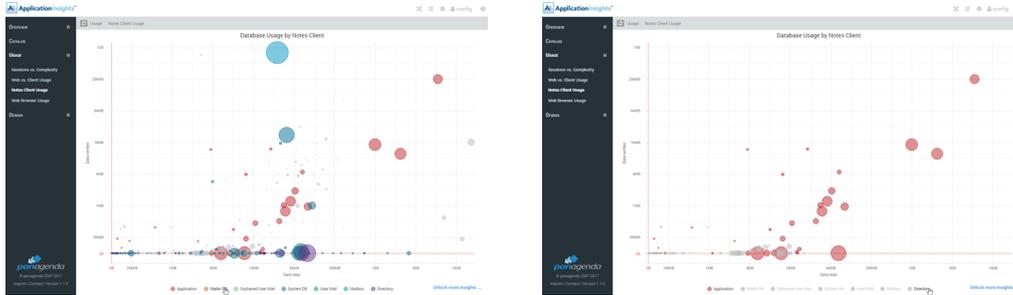


Figure 2: Subset Selection

Certain Bubble charts allow you to focus in on part of the graph by using the mouse to select the top left most grid point of the area you would like to focus in on, holding the right mouse button down and drawing it to the bottom right grid point you would like to focus in on. Releasing the mouse there. This will concentrate the graph in on the selected area.

Example views where this is used: Usage - Sessions vs Complexity, Usage – Web vs. Client Usage, Usage – Notes Client Usage, etc.

3.4.2 Maximize screen usage

Hide left menu: To maximize screen width on a graph or page it is possible to hide the left-hand menu of panagenda ApplicationInsights. To do so click the so called 'hamburger' icon in the top right corner.

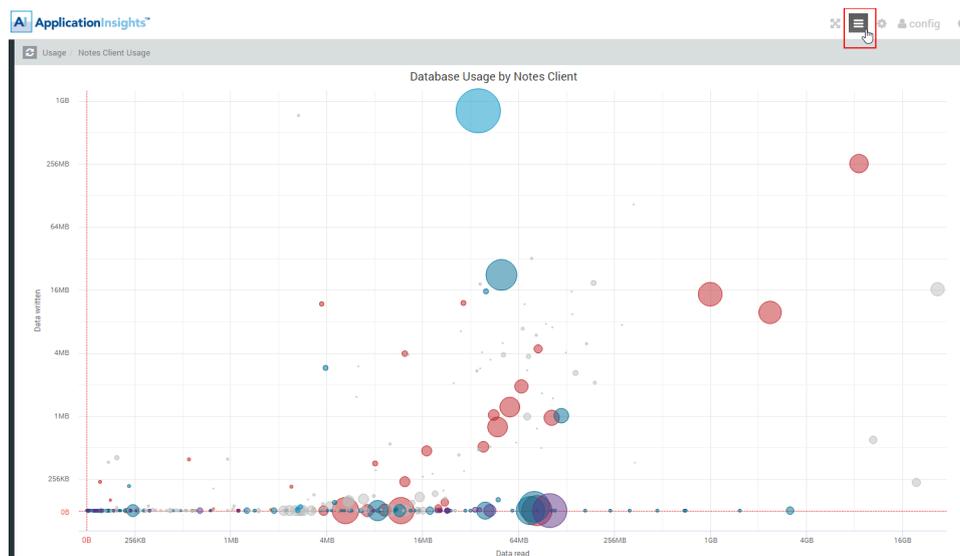


Figure 3: Collapsed Menu

To return the menu simply click the same icon again.

3.4.3 Help and additional information

Certain elements in the panagenda ApplicationInsights interface have Help text associated with them to explain the specific graph or element.

Underlined: An underlined (dashed line) element usually indicates that on mouse over a help text or additional information is available.

“?” Symbol in right corner next to graphs: Indicates that there is a help text available that will be visible on mouse over.

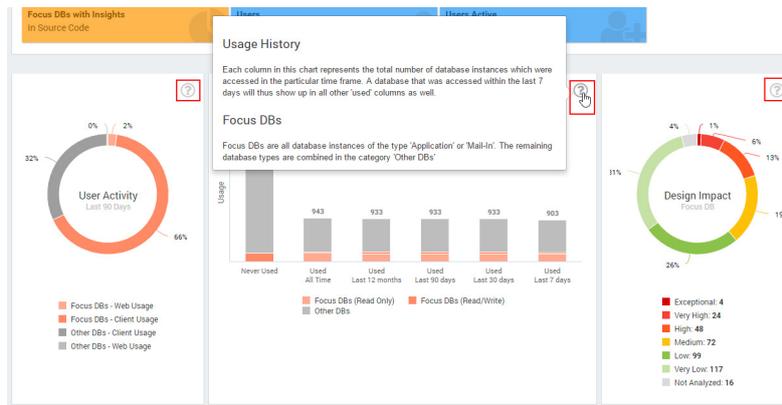


Figure 4: Help Symbol

Details: Some graphs will allow you to mouse over or click on a node or data point to access more data. To go to the instance details page for any given database in any of the bubble charts for instance, simply double click on the bubble.

4 AI Components

This chapter focuses in on the individual graphs and overviews panagenda ApplicationInsights offers you.

4.1 Overview

The overview section of panagenda ApplicationInsights gives you a quick overview of your complete application environment and the status of data collection.

4.1.1 Status Dashboard

The status Dashboard gives you a first glance of the usage and design insights status:

- **# DB Instances Discovered:** The rounded number of database instances (individual database replica's) discovered on the servers being polled.



Note: *If not all servers in the environment are included in the collection this number might be skewed. For accurate information, it is important to include at least all servers that might hold applications or mail-in databases but preferably all servers for proper visualizations across the entire Domino environment*

- **% Databases Used in Last 90 Days:** The percentage of database instances (out of the discovered instances) accessed by users in the last 90 days
- **# Focus DBs Identified:** The number of database instances identified as either a type "Mail-In" database or type "Application". Other databases are: user, server & orphaned mailboxes, system databases and Directories.
- **% Focus DB Instances in License Scope:** Percentage of database instances considered "Focus databases" that are covered for analysis by your license agreement
- **% Focus DBs Used in Last 90 Days:** The percentage of focus database instances (out of the discovered instances) accessed by users in the last 90 days
- **# Focus DBs Analyzed:** The number of Focus databases for which design analysis was successfully performed. If this number is lower than the number of Focus DB's identified this could indicate that not all databases could be analyzed.

- **% Focus DBs with Design Insights:** Percentage of the database instances analyzed for which design Insights were found. Design insights consist of “Findings” of specific code elements or combinations of code elements, based on predetermined rulesets.
- **# Users Registered:** Total number of users registered in the Domino environment (taken from your Directory infrastructure)
- **# Users Active in Last 90 Days:** Number of users who have shown activity in the Domino environment over the last 90 days.

User Activity Last 90 Days: Shows you the activity generated for focus versus non-focus databases based on the “user access days”. A user access day is the equivalent of a user accessing the database on a given day. By using days instead of sessions, it prevents sudden spikes of activity skewing the data. This graph gives you insights into which types of databases most of your users are working in (focus versus non-focus databases) and how they are accessing the databases (web versus client).

Usage history: Each column in this chart represents the total number of database instances which were accessed in the particular time frame. A database instance that was accessed in the last 7 days will thus also show up in all other ‘used’ columns. This graph will of course get more meaningful over time as it will start to highlight how many are still in daily use and how many are silently falling into disuse.

Design Impact: Indicates how the focus databases (databases of type Applications or Mail-In databases) are categorized in terms of design complexity. Design complexity depends on many factors which will be further explained in “Design Complexity” on page 35.

4.1.2 Environment Overview

The Environment overview provides you with a single overview of how the infrastructure of your Domino Environment is set up.

It lists information about the number of:

- Directories
- Servers
- Servers analyzed (by panagenda ApplicationInsights)
- Number of Domino Releases (utilized by the servers)

- Number of Operating Servers (utilized by the servers)
- Registered Users
- Active Users
- Groups
- Group members
- ACL entries
- Databases deployed (focus & non-focus databases)
- Database storage used
- Databases touched by active users
- Views
- Views indexed
- View storage

4.1.3 Inventory

The Inventory page will give you a global overview over the total number of database instances, and replica sets you have.

By comparing these numbers against usage information and disk space used you get a better insight into the actual capacity utilized by your Domino environment.

4.2 Catalog

The Catalog section of panagenda ApplicationInsights is where you can get deep dive information about a specific database instance. Both in terms of usage as well as design.

4.2.1 Working with the Catalog List

Instances listed:

The catalog list lists all instances your license entitles you to access for analysis. Please review the paragraph on licensing to understand how the selection of instances you see is made. As this list displays instances, it is possible that you see several instances of the same replica set.

Quelle der Admins	cronus/panagenda	panagenda/adminquelle.nsf	■ Application	-	2017-04-07	101	Low
Analyze PREV	cronus/panagenda	panagenda/analyze_prev.nsf	■ Application	-	2017-04-08	115	High
Contacts	mar duk/panagenda	crm/contacts.nsf	■ Application	900-00-00 7.2 SR1 DE	2017-04-08	1	Exceptional

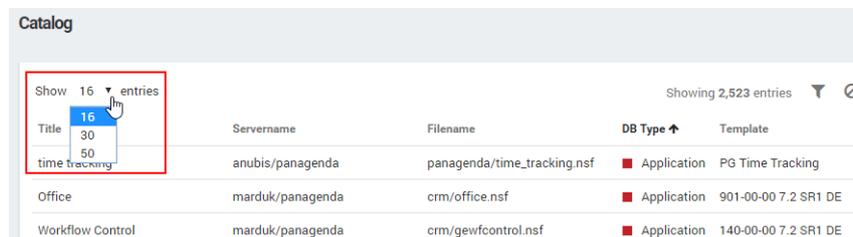
← 1 2 3 ... 158 →

Notice: 583 applications omitted from catalog due to licensing restrictions

Figure 5: Instances listed

Displayed number of lines:

The “Show __ entries” option in the top left allows you to show more or less values per page.



Catalog

Show 16 entries

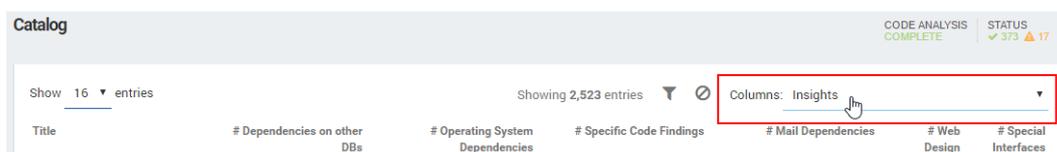
Showing 2,523 entries

Title	Servername	Filename	DB Type	Template
time tracking	anubis/panagenda	panagenda/time_tracking.nsf	■ Application	PG Time Tracking
Office	mar duk/panagenda	crm/office.nsf	■ Application	901-00-00 7.2 SR1 DE
Workflow Control	mar duk/panagenda	crm/gewfcontrol.nsf	■ Application	140-00-00 7.2 SR1 DE

Figure 6: Show x entities

Displayed columns:

The “Columns” option in the top right allows you to select one of the column sets. The “Default” set contains information like the location, the type of instance and the impact score.



Catalog

CODE ANALYSIS COMPLETE STATUS 373 17

Show 16 entries

Showing 2,523 entries

Columns: Insights

Title	# Dependencies on other DBs	# Operating System Dependencies	# Specific Code Findings	# Mail Dependencies	# Web Design	# Special Interfaces
-------	-----------------------------	---------------------------------	--------------------------	---------------------	--------------	----------------------

Figure 7: Displayed Columns

The “Usage” set show information about the number of sessions and the types of access (mobile, web or Notes client) as well as some design information.

The “Scores” set shows you, among others, information about ranking and scores for complexity, insights and impact.

The “Insights” sets give you specific filter options to select databases with similar insight findings. Both on a ruleset (category) as well as Finding level.

Sorting & filtering:

To find a specific instance you can sort the list on various elements. For example, on the name of the database, the server the instance is located on, or even on a type or impact score if you are interested in finding databases of a specific type or complexity. To sort simply click on the column header.

Alternatively, you can also filter the list on any of the columns by clicking the “funnel” icon in the top. The column headers will then become fields where you can enter the search term or value you want to filter on.



Figure 8: Sorting & Filtering 1

To clear a set filter use the “stop sign” icon next to the funnel icon.



Figure 9: Sorting & Filtering 2

Opening a database instances details:

To open the details of any instance simply click on the table row for the instance you are interested in.

4.2.2 Instance Details - General

On opening the database instance details, you will be presented with a header that contains general information about this specific instance. Keep in mind that if a database has multiple instance in its replica set each of them will be analyzed separately and classified separately. This is done deliberately to ensure that you get the most accurate information.

A database might for instance have replicas on three servers but all usage might only be taking place on one server. Also, due to replication restrictions or problems, it is possible that the design on databases sharing the same replica id is not the same.

Therefore, the usage, design and ranking information portrayed in the top is all instance based.

4.2.3 Instance Details - Tab Application

The Application tab is divided into three sub tabs (left side)

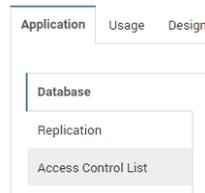


Figure 10: Application Tabs

Database:

Displays general information about the database instance as taken from the database properties, the size of the database instance, the template it inherits from (or is) and whether it is listed in the Domino catalog.

Replication:

Indicates whether replication is enabled and which settings have been activated

Access Control List:

Gives information about how many entries there are in the ACL with a specific level of access and which level of access special user entries like "Default" and "Anonymous" have.

4.2.4 Instance Details – Tab *Other Instances* (x)

This tab is only visible if other database instances have been identified that share the same replica id. If there is no tab for instances then no other replicas were found on the analyzed servers. The value between brackets behind the tab name indicates how many other instances share the same replica id.

Each instance will be identified with its own tab showing that instances "Application" details and a button to open the full analysis details for that instance.

Keep in mind that multiple instances with the same replica id can exist on one server.

4.2.5 Instance Details – Tab Usage

The Usage tab gives you specific information about how the database instance is being used. It shows you the number of 'user days' activity is taking place in the instance, the number of active users per period, the amount of data being uploaded/written versus downloaded/read and the number of web and client sessions.

Having this insight can help you determine whether a specific instance needs to remain as a separate replica or could be consolidated / removed.

4.2.6 Instance Details – Tab Design Complexity

This tab is only visible if a design analysis was performed on the database and if it is in scope.

The design complexity tab gives insight into the design elements of the database. Giving you the totals for each type of element being used, how many of them contain code – and what type of code as well subsidiary elements like fields, subforms, etc. are being used.

LOC stands for Lines of Code

Ø for average

for number

4.2.7 Instance Details – Tab Design Similarity

The Design Similarity tab is where you will be able to identify databases with near similar designs. Regardless of whether they share a template or not.

If an instance belongs to a similarity cluster (meaning it shares a design that is for 95% or more similar to that of other instances in the environment) this tab will appear.

On the Similarity tab each node in the diagram represents a sub-cluster consisting of 1 or more instances which share the exact same design. The top node represents the design held by the instance whose details you are currently reviewing. The other nodes represents sets of instances with near similar designs. By clicking on an individual node, you get to see additional info about the instances sharing that design. Like the names and templates each instance inherits its design from (if applicable). By holding the CTRL key and selecting multiple

nodes you can then identify which other database are sharing a near similar design and could potentially be candidates for template consolidation.

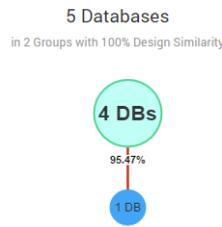


Figure 11: Design Similarity

4.2.8 Instance Details – Tab *Design Insights*

This tab is only visible if a design analysis was performed on the database, if design findings were found and if it is in scope. In all other cases the tab will not be visible

This tab is where panagenda ApplicationInsights helps you identify potential problems and challenges in the code of your database instance design.

Based on a predetermined list of criteria (rulesets) panagenda ApplicationInsights identifies certain occurrences or combinations of occurrences in your code that could require attention in case of optimization, modernization, transformation or migration.

Findings:

Findings are categorized into ruleset categories like:

- Web Design
- Mail Dependencies
- Special Interfaces
- Operating System Dependencies
- Dependencies On Other Databases
- Specific Code Findings

Each ruleset can be opened (double click) and will list the type of findings in that category and the number of occurrences in the code. Each finding criteria has a weight that help identify the potential risk of the finding. This is indicated by the color of the finding label.

Findings classified as more high risk/more complex are shown as red while less important findings would be indicated with green.

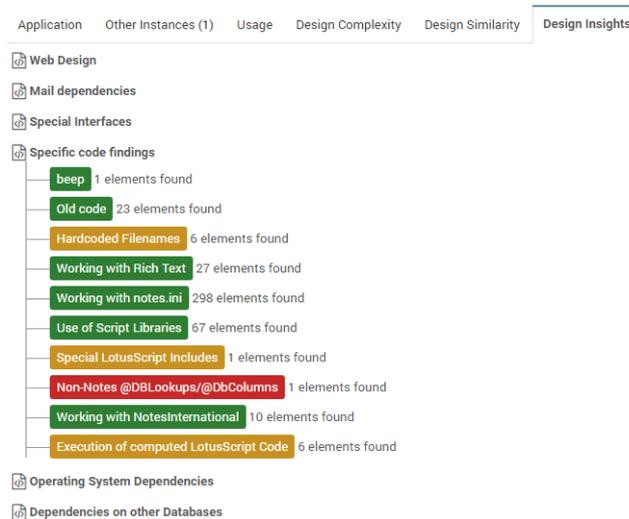


Figure 12: Specific Code Findings

Each finding category can be opened to review the individual code snippets and their location.

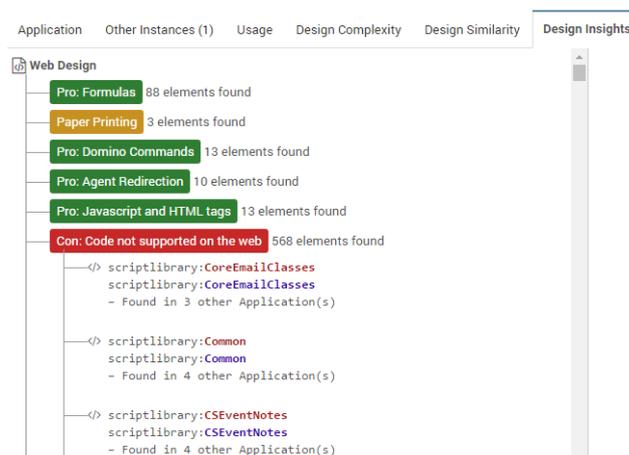


Figure 13: Web Design

Each finding lists the design element type, design element name and the action in which the code is located. It also lists in how many other applications that code element can be found.

Not only does panagenda AI give you the indication of what code elements could be problematic, it also tells you in how many other applications the same code is being used. Allowing you to plan more efficiently and replace code across your platform as you go.

To view the affected code simply click on the element block. The code will appear in the right hand window.

The affected code lines are highlighted in yellow and their location is visible in the right scroll bar for easy navigation. Use the “<< previous” and “next >>” buttons in the top of the code window to easily navigate tot the affected code lines.



```

Design Insights
<< Previous Next >>
1 xxTeamMembers:=@Unique(@Trim(@GetProfileField("calendarprofile";"ReadCalendar")):@GetProfileField("calendarpro
2
3 REM {@If(@Trim(xxTeamMembers)="";@Prompt([Ok];"IBM Lotus Notes";"There are no Team Members specified yet.") &
4
5 @If(@Trim(xxTeamMembers)="";@Prompt([Ok];"IBM Lotus Notes";"There are no Team Members specified yet.") & @Retu
6
7 @For(index := 1; index <= @Elements(xxTeamMembers); index := index + 1;
8   xxMemberName := xxTeamMembers[index];
9   @If(@Contains(xxMemberName;"");xxMemberName="";xxMemberName);
10  @If(@Contains(xxMemberName;"-Default-");xxMemberName="";xxMemberName);
11  xxTeamMemberList := @Trim(@If(@Elements(xxTeamMemberList); xxTeamMemberList : xxMemberName; xxMemberN
12
13 @SetField("EnterSendTo";@Unique(@Trim(EnterSendTo):@Trim(xxTeamMemberList)));
14 @SetField("AddTeamMembers";"1");
15 @Command([RefreshHideFormulas]);
16 @Command([ViewRefreshFields]);
17 @Command([EditGotoField]; "EnterSendTo");
18 @Command([EditInsertText]; "");
19 @PostedCommand([ViewRefreshFields])
  
```

Figure 14: Affected Code Lines

4.3 Usage

The Usage section provides insights into the clients used by users to access the various database types. These overviews will help you determine which databases are predominantly used in the IBM Notes client and which predominantly through web browsers.

By linking the type of client used to details about the actual usage and number of users accessing that database, panagenda ApplicationInsights allows you to prioritize and determine importance and business value of each database.

4.3.1 Sessions versus Complexity

The Sessions versus complexity view allows you to see which of your focus databases instances with a high complexity are also often accessed. Each bubble represents a single database instance. The size of the bubble indicates how many unique users accessed the database over the analyzed period. In the details window that becomes visible on moving the mouse over the bubbles you find additional info on the type of clients and the “Access days” (see “Definitions” on page 11). To see full catalog details on the specific database instance simply double click the bubble to go to the instance details.

The closer a bubble is to the top of the scale, the higher the complexity score. The more to the right a bubble is the higher the amount of sessions. Database instances to the left bottom corner are in general easier targets for archiving as they are less frequently used and less complex. The same goes for smaller bubbles which indicates less users are involved.

Grayed out bubbles indicate databases out of scope of the license key obtained by your organization.

4.3.2 Web versus Client Usage

In this overview, all databases that show activity are represented as bubbles. The graph uses various elements that help you identify information about the databases in terms of number of users, type of access and activity.

Positioning on the graph:

The top dashed line represents pure web access, meaning that the databases closer to the top are predominantly or exclusively accessed through browsers.

The bottom dashed line represents pure IBM Notes client access. Meaning that the databases are predominantly or exclusively accessed through IBM Notes Clients ("rich" clients).

'Mixed access' indicates those databases that are accessed through both rich clients as well as browsers.

The horizontal position of the bubble indicates the number of sessions (web & client) generated by users accessing the application. The further to the right a database is the more activity a database has had.

Color & Size of the bubbles:

Each bubble represents a database instance.

The size of the bubble indicates the number of unique users having accessed that database in the period analyzed.

The color indicates the type of database. The legend in the bottom explains which color represents which database type. Click on a type in the legend to remove that type from the graph to focus on for instance just applications.

Hover over / details:

To see the details for any specific database simply hover over the bubble. Note that for non-licensed instances, less info is displayed or info is obfuscated. To see full details simply double click the bubble.

The details popup shows:

- Name, location and replica id of the database
- Complexity score or "-" if no score is awarded.
- Number read/download versus write/upload actions in the analyzed period
- Number of sessions (used for the vertical and horizontal position on the graph).
- Number of unique users (used for the size of the bubbles)
- User Access Days per type of session

4.3.3 Notes Client Usage

On the notes client usage tab, you can see the type of activity being performed on databases accessed through the IBM Notes client.

As in the previous graph the size and color of a bubble indicate the type and number of users accessing the database. The horizontal position indicates the amount of data traffic caused by READ operations performed on this database and the vertical position the amount of data traffic caused by WRITE actions.

The more a bubble is towards the left bottom corner of the graph the less used it is in rich clients. Similarly, the more towards the right upper corner the more actively used the database is on rich clients. Whereby the size of the bubble indicates the number of unique users performing these actions.



Note: A database showing as having little to no activity in this view does not necessarily mean it is not being used. It could be that it is predominantly used through web clients, for instance.

4.3.4 Web Browser Usage

Displays data uploaded (WRITE operations) versus data downloaded (READ) operations for databases accessed through web browsers. See “Notes Client Usage” on page 30 for information on how to use this view.

4.4 Design

The Design section of panagenda ApplicationInsights gives you global information about the design complexity and details of your application landscape.

4.4.1 Complexity

The Complexity page lists the overall totals of all analyzed designs in your application landscape when it comes to focus databases.

- Number of focus databases analyzed
- Total number of lines of code – use the mouse-over to see a breakdown in terms of used coding language
- Average complexity index
- Highest complexity index

The average & highest complexity index give you an idea of the spread of complexity. By looking at the efforts and costs of transformation or modernization of a low complexity, an average complexity and a high complexity database you can get an impression of time and effort needed to perform the actions you can give yourself a starting point to do calculations for planning and investment. Combining this with other information (like usage) allows you to extrapolate that over your entire environment.

Additionally, this page gives you insights into the prevalence of the various coding languages used, the number of design elements and the number of code lines represented by those design elements.

At the bottom of the page you see a distribution overview of all databases with a complexity index of High, Very High or Exceptional.

4.4.2 Similarity

The Similarity tab is where you will be able to identify database that share a design but not necessarily a template, or designs and/or templates that are so similar that it could be beneficial for optimization and modernization to consolidate them into one single template.

To do so, panagenda ApplicationInsights analyzes the designs of each of your focus databases and compares them. Both on design elements as well as on individual code line level. Databases that share at least 95% or more similarity with each other are then grouped into 'similarity groups'. Independently of whether they share a design template or not.

Overall Similarity insights:

The top bar of the Similarity page shows indicators for the number of similarity clusters that have been identified, how many instances would be impacted if all these similarity clusters were indeed consolidated into templates and how many instances are included in the biggest cluster. A first indicator of the potential impact of template consolidation.

Similarity Cluster distribution:

Below that a graph will show you the various identified similarity clusters. Larger bubbles represent larger clusters (more database instances sharing the same similarity) and smaller bubbles represent situations where only two or a few database instances are involved.

The bubble color is determined by the average similarity score inside a cluster. The higher the similarity the darker the bubble. The exact percentage is also displayed as text if bubble size allows.

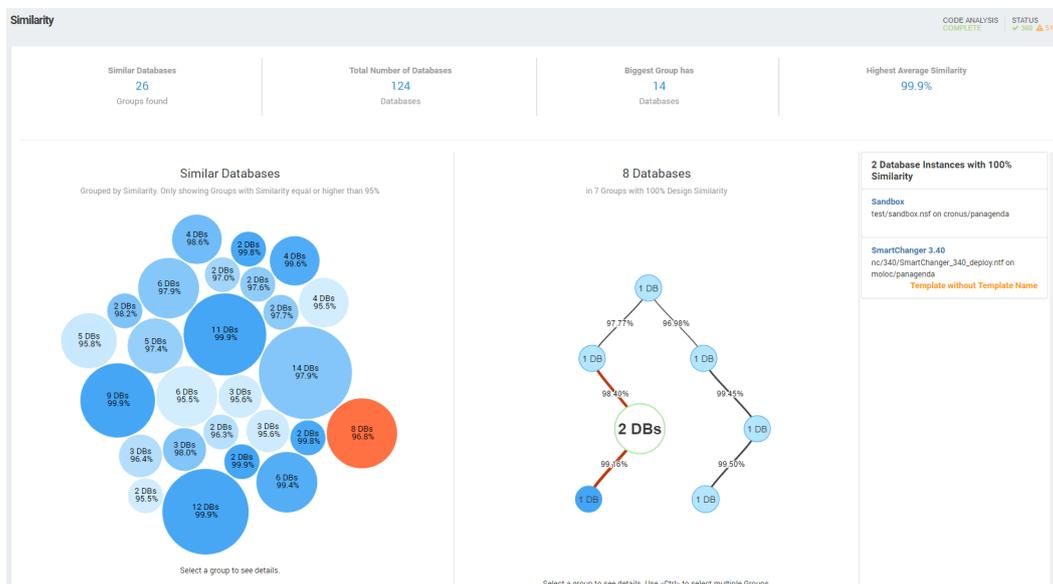


Figure 15: Design Similarity

Similarity Cluster details:

By clicking on a cluster bubble a second column will appear holding a node diagram representing that cluster's structure. Each node in this diagram represents a sub-cluster of 1-n database instances that share a unique design with 100% similarity score, with the size and number on the node representing how many instances share that same unique design. Links between nodes show the percentage of design similarity between sub clusters.

Instance and template identification:

To get more information about the database instances represented, simply click on a node. A third column of information will then appear which is a list of instances represented by the selected node. For each instance the title, location and template name are shown so that you can quickly identify if these instances indeed inherit from a shared design or perhaps should do so.

You can select more than one node at a time by holding the CTRL key on your keyboard. To see an individual instances full details and design insights simply click on its title in the third column.

By using this page, you should be able to quickly identify which database designs within your environment can be consolidated and optimized, which databases should be assigned to a template and what the impact will be when performing template consolidation.

4.4.3 Insights

The Insights page gives you information about the designs found containing insights ("Findings") based on the rulesets. Here too you see information about the number of lines of code (and on mouse-over their distribution across various coding languages) but you also see the number of databases that have findings and information about the total number of findings found.

A finding is where a specific code element or combination of code elements was found that is listed as something that requires consideration in case of for instance modernization. Think of things like dependencies on other databases, operating system or rich client but also old code, code not supported on the web and mail dependencies.

These findings are categorized into the following rulesets:

- Web Design
- Mail Dependencies
- Special Interfaces
- Operating System Dependencies
- Dependencies On Other Databases
- Specific Code Findings

All these various categories of findings are then listed in the “Design Insights – by Ruleset and Findings” graph. To focus in on a certain ruleset simply click on it to expand it for better overview. To collapse it back to the overview simply click on the button in the right upper corner.

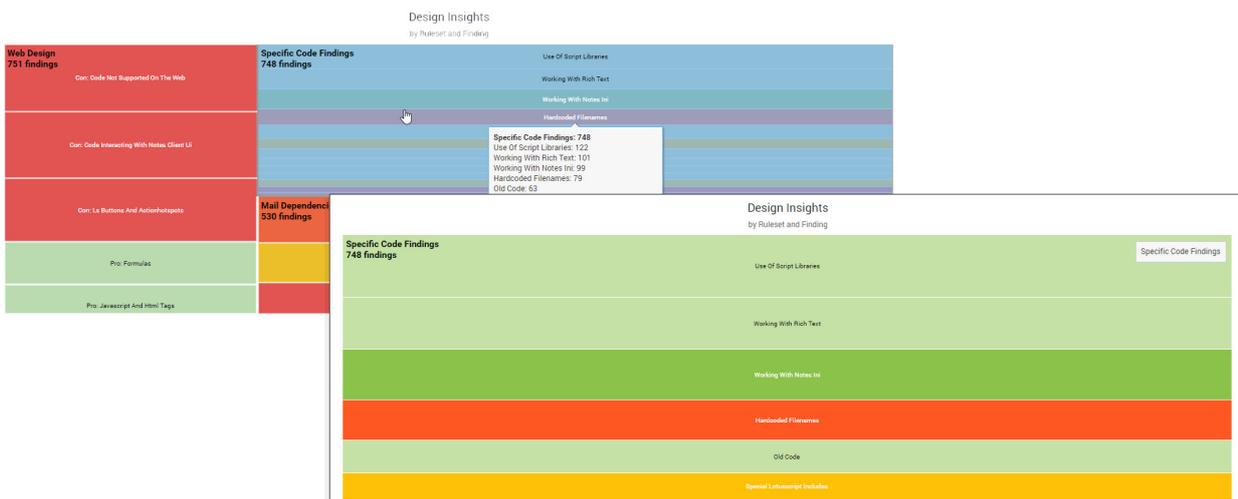


Figure 16: Design Insights Ruleset

The color of each finding indicates its relative impact (Dark red being severe and green being low).

The distribution of the 6 ruleset categories as well as the findings within them give you an insight into the type of potential roadblocks or focus areas in case of design changes you might face. To see which databases are within each ruleset use the Catalog view and sort them based on the various “Insights” column sets.

5 Design Complexity

panagenda ApplicationInsights defines the design complexity of a database on several factors and criteria. The idea being that it should help you determine how to best address topics like modernization and migration.

By combining the design complexity with the actual application usage, a strategy can be determined that is realistic and cost effective.

5.1 Collected Design Data

panagenda ApplicationInsights analyses a database design based on:

- Design elements: Identified as Domino design elements like forms, views, agents, script libraries, actions, fields, etc.
- Code blocks: A code block is a set of code lines that together form a logical code element. Like for instance the formula code of an action within a form.
- Code lines: Individual lines of code

For each of these elements various information is collected and analyzed. Like for instance the design element type, the code language, the root element and relative path of the analyzed element and the actual code itself.

By doing so panagenda ApplicationInsights can give you deep dive information about complexity and similarity of database designs.

5.2 Design Complexity and Impact Criteria

Based on the collected data panagenda ApplicationInsights uses a wide range of criteria to determine design complexity and impact.

Factors that for instance can influence the complexity and impact are:

- Type and number of specific design elements occurring in the application design
- Specific combinations of design elements occurring in the application design
- Hardcoded references (for instance hard coded server names, etc.)
- Number of code blocks and code lines in the application
- Domino/Notes client dependencies
- Calls to local systems (for instance code that references a user's local hard disk or accesses specific files stored on the server)
- Occurrence of specific code elements that can cause potential migration/modernization challenges
- Occurrences of Author/Reader fields
- SMTP / mail / fax / etc. references that indicate a dependency on a specific system or server
- Third party software dependencies
- Encryption dependencies
- Mail dependencies
- Occurrences of specific code functions and snippets or combinations of such code functions and snippets.
- Etc. etc.

5.3 Weighted analysis

An important part of the design analysis is that criteria are not just analyzed on their own but also used in combination with others. Meaning that the design analysis is based on weighted criteria, not just individual elements. What does this mean? Well let's say an application sends out a mail when a user performs a certain action. In the code this is identified by, for example, a formula like @mailsend. However, if that mail is also encrypted in the same code block before being send out the complexity is much higher as this means that it hits at least three different criteria:

- a) It uses mail functionality

- a) It uses encryption
- a) It has a Notes client dependency (for the encryption to work)

The total weighted complexity of the combination of all three is higher than the sum of each individual element would be and therefore panagenda ApplicationInsights will give it a higher complexity rating.

5.4 Design Complexity Ruleset

To be able to give insight into the complexity of an application, weighted rulesets are defined for panagenda ApplicationInsights. These are based on our experience with analyzing database designs for modernization.

The following six rulesets are provided:

- Web_design
- Mail dependencies
- Special Interfaces
- Operating System Dependencies
- Dependencies on other databases
- Specific code findings

If your organization requires additional custom rulesets, please contact panagenda.

6 Support & Help

Additional details are available in technote 2001728 - Q&A about panagenda ApplicationInsights and the offering for entitled IBM customers with link:

<https://ibm.co/2plriSL>

If additional support is needed, please contact support or submit a ticket. To do so click in the right upper corner on the "?" icon and follow the links "**Support**" or "**Submit a ticket**".

DISCLAIMER

panagenda, panagenda product names and all related logos are trademarks owned by panagenda. All other names of products and enterprises in this documentation are the property of their respective owners.

panagenda reserves the right to update this documentation without being obliged to announce the changes or revisions.

Although all due care has been taken in the preparation and presentation of this documentation, the corresponding software may have changed in the meantime. panagenda therefore disclaims all warranties and liability for the accurateness, completeness, and currentness of the information published, except in the case of intention or gross negligence on the part of panagenda or where liability arises due to binding legal provisions.

Limitation of liability for external links

This documentation contains links to the websites of third parties ("external links"). As the content of these websites is not controlled by panagenda, we cannot assume any liability for such external content. In all cases, the provider of information of the linked websites is liable for the content and accuracy of the information provided. At the point in time when the links were placed, no infringements of the law were recognizable to us. As soon as an infringement of the law becomes known to us, we will immediately remove the link in question.

NOTES