



G Data Business Solutions

Reference Guide

G Data Application Development

Table of Contents

Introduction	5
Section A: Planning and deployment	6
1. Network and client management.....	6
1.1. Network layout.....	6
1.2. Security components	11
2. Choosing a solution	12
2.1. G Data business solutions.....	12
2.2. System requirements	13
2.3. Licensing.....	15
3. Installation scenarios	16
3.1. Local deployment.....	16
3.2. ManagedService.....	21
4. Deployment	23
4.1. Preparation	23
4.2. Clean installation	24
4.3. Upgrade installation.....	27
4.4. Initial configuration	27
4.5. Server updates and registration.....	29
4.6. Server database backup and restore	30
4.7. Client deployment.....	32
4.8. Finalizing deployment.....	38
5. Remote administration	39
5.1. Desktop application.....	39
5.2. Browser.....	40
5.3. Mobile.....	42
5.4. MasterAdmin.....	43
Section B: Using G Data business solutions	44
6. Dashboard and monitoring	44
6.1. Statistics	44
6.2. Reports	46
6.3. ReportManager	48

7. Managing clients	50
7.1. Using groups	50
7.2. Integrating Active Directory	51
7.3. Signature and program file updates.....	51
7.4. Removing a client.....	55
7.5. Performance	55
8. Real time protection	57
8.1. Traffic scans.....	58
8.2. File system.....	59
8.3. End user security permissions.....	62
8.4. Performance	63
8.5. Operating system security.....	64
9. On demand protection	67
9.1. Idle scan	67
9.2. Scan jobs.....	68
9.3. Exceptions.....	74
10. Handling a malware infection.....	75
10.1. Automated detection and mitigation.....	75
10.2. Extended mitigation.....	77
10.3. Analysis.....	78
11. Mobile device management	80
11.1. Managing mobile devices	80
11.2. Real time and on demand protection	80
11.3. Device policies	81
11.4. Apps.....	82
11.5. Contact management and filtering	83
11.6. Anti-theft	84
12. Backups	86
12.1. Managing backups.....	87
12.2. Create a backup.....	88
12.3. Restore a backup.....	92
13. Firewall	94
13.1. Managing firewall clients.....	94
13.2. Autopilot.....	96

13.3. Rule sets	96
13.4. End user permissions	98
13.5. Logs	99
14. PolicyManager	101
14.1. Applications.....	101
14.2. Devices	103
14.3. Web content.....	105
14.4. Internet usage time.....	106
15. Patch management	108
15.1. Patch management cycle	108
15.2. Standardized patching policy	110
15.3. PatchManager.....	110
16. Mail server security.....	116
16.1. Exchange plugin	116
16.2. MailGateway.....	118
17. Advanced configuration	137
17.1. Configuration files.....	137
17.2. Configuration tools.....	142
17.3. Debug logs.....	146
Acronyms	150
Index	151

Introduction

G Data provides high-end malware protection for SMB and enterprise networks. The solutions are based on central configuration and administration and provide as much automation as possible, while allowing extensive customization. All clients, whether workstations, notebooks, file servers, or mobile clients, are administered centrally. Client processes run invisibly in the background and automatic internet updates enable extremely fast response times. G Data supports various administration approaches, whether security functions need to be configured to function autonomously, or full control over the software's actions needs to be maintained. This document will support informed decision making about the deployment of G Data business solutions and will provide recommendations and optimal settings to protect SMB and enterprise networks.

The process of providing network and client security can be divided in three parts. Ideally, security is kept in mind right from the start, before deploying any hardware or software. However, even existing networks can and should profit from a well-wrought security policy. The planning stage helps administrators think about the needs and wishes of the end user, the physical possibilities of hardware, the optimal layout of the enterprise network, and the layers of security that will be added to all nodes. With a high-level network layout in mind, an informed decision can be made about the G Data security solution that will be deployed.

With the basics arranged, administrators can move on to the actual deployment of their G Data business solution of choice. Whether it is an SMB network with fewer than 50 clients, or an enterprise deployment with over 1000 clients, the installation of G Data software can be streamlined and adapted to each situation. To enable swift client deployment, G Data solutions support several client installation scenarios, ranging from automated remote installations using Active Directory to local client deployment.

After deploying server(s) and clients, the final phase is initiated. Using the newly built client-server infrastructure, administrators can configure client protection, as well as mail server security, backups, patch management, security policies, and much more. This guide will assist in effectively configuring G Data software for optimal security without sacrificing performance.

Section A: Planning and deployment

1. Network and client management

Setting up and managing a secure business network can be challenging. Network and client hardware and software need to be configured to support various end user workflows while keeping out unauthorized users, attackers and other threats. Rather than immediately deploying G Data software to all existing servers and clients in the network, network layout and client management should be considered first. By organizing the network into different zones and defining client roles, subsequent configuration can be made significantly easier. By dividing the network and using standardized client profiles, time is saved when deploying new security updates, planning scans or scheduling backups. Additionally, critical parts of the network and their clients can be defined as such, allowing for focus on the most important infrastructure when necessary.

1.1. Network layout

The network layout concerns the physical arrangement of all network hardware, such as modems, routers, switches, servers, clients and other networked devices. Rather than adding networked devices as they are acquired and deployed, working with a standardized network layout allows network administrators to maintain an overview of the network as a whole. By working with network zones and client roles, a standardized configuration can be deployed to every new device as soon as it is added, saving time and ensuring compliance across the whole network. This concept applies to smaller networks as well as bigger ones. As soon as more than one client is used in a business workflow context, having a default configuration prevents frustration when troubleshooting client issues, security problems or network irregularities.

Building a network can be planned by creating a network diagram. The diagram should list all network devices, including routers, switches and other supporting equipment. It provides an overview of the various physical security layers, including modems, routers, and firewalls. Before deploying any software security solution, these devices, serving as the network's first line of defense, should be configured. For example, when using a modem or router with built-in firewall, make sure that the firewall is enabled and that appropriate rule sets have been defined to drop malicious traffic. Other built-in security functions should be assessed and if appropriate, enabled. Note that all of these settings are to be considered a baseline of network security: enabling a firewall alone does not protect network clients. Once a security solution has been deployed to client devices, previously configured measures on hardware devices may have to be reconfigured to ensure compatibility.

In addition to visualizing the various hardware security layers, a network diagram is a very productive way to group clients together. This is where the concepts of network zones and client roles come in. In its most basic form, a network zone is a specific segment of the network which has been assigned for a specific purpose. Network zones allow security measures to be configured per IP range, giving administrators immediate insight in the required policies for a new device deployment within a specific zone. Grouping by network zone means grouping devices by parameters such as physical location, purpose, security restrictions or any other properties. For example, devices can be grouped into a network zone based on their physical location in the building ("Sales department", "Front office"). When dividing devices based on security restrictions, a network zone could include internet-facing servers (DMZ), restricted-access local clients, or any other device group that should be configured using common policies. Especially the deployment of a DMZ is recommended for

organizations that host their own internet-facing servers (such as hosting a website, a mail server or an FTP server). Configuring these machines as a separate network zone, with strict firewall rules separating them from the internal network, hugely reduces the risk of an attacker successfully infiltrating devices in the internal network. Each network zone can be seen as its own trust zone with its own security restrictions. Traffic flow to and from each network zone can be restricted on a network level to make sure that critical infrastructure cannot be accessed without authorization. Even for existing networks, defining at least a few different trust zones can help get an insight which machines are most essential to the company and should be protected accordingly.

Every network client can be assigned a specific client role, corresponding to its use, its priority, its security risk or other parameters. For existing networks that have their clients organized in Active Directory (AD), a client role roughly corresponds to an AD Organizational Unit. A client that will be used mostly for basic office tasks could be assigned the role "Office". For priority-based roles, a generically deployed client for office activity could be assigned a lower priority client role than the meticulously configured developer's client. The client role defines its local security policies as well as its software deployment. There can be some overlap between the client role and its network zone, for example if network zones have been configured based on physical proximity. Ideally, there will be as much overlap as possible, so that administrators can configure one policy per zone that covers all its clients. For example, the proximity-based network zone "Sales department" would ideally only contain devices in the "Sales" client role, which would allow administrators to deploy the same software configuration and security policies to all clients in that network zone. However, not all clients within one network zone might be used for the same purposes. When using a network zone that is based on security permissions, the clients will share a large part of their security settings, but can still have different software deployments and thus different client roles. The following diagrams will clarify the concepts of network zones and client roles in several common network deployments.

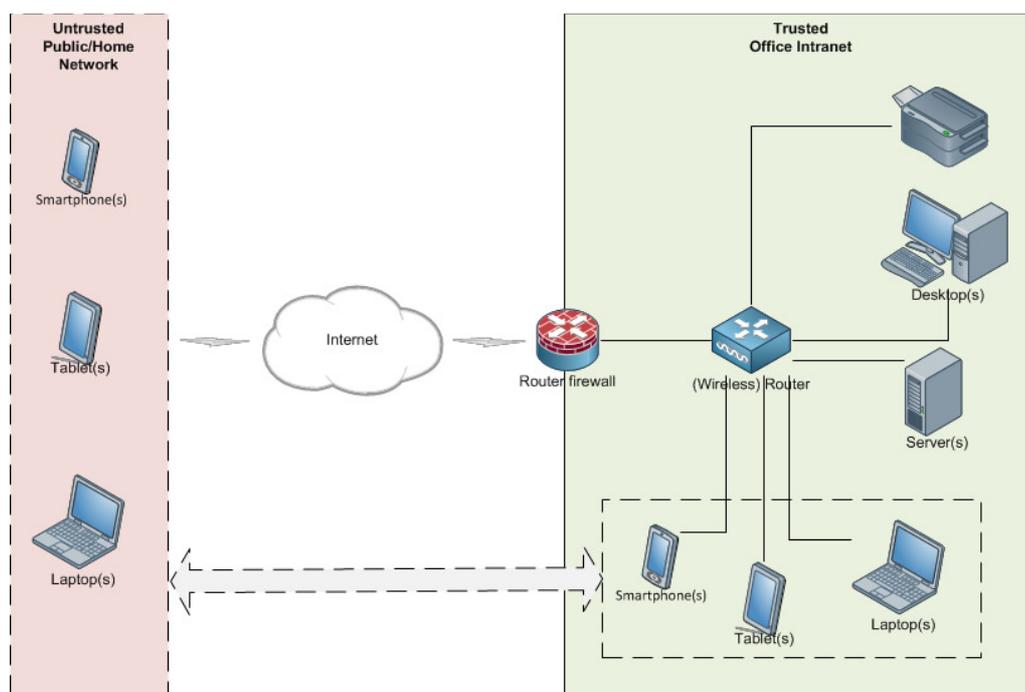


Image 1: Small office network

The small office network strongly resembles a home deployment and is typically deployed in networks that have only one server and relatively few clients. Schematically, it is the simplest network layout. All network clients are located in the same network zone and connect directly to the

internet via a (wireless) modem router. Hardware-level network security is provided by a built-in router firewall. The clients may or may not have been assigned different client roles. Some clients, such as smartphones, tablets, and laptops, are used within the trusted office network, as well as on untrusted public or home networks. They are not only managed within the office network, but should also be appropriately secured when they are being used in untrusted contexts to connect to the office network.

In the small office network scenario, most security measures need to be implemented on the clients. Other than the optional hardware firewall, no network-level security is present. There is no clear distinction in network zones, but client roles should still be enforced, even though there may only be one or two machines per role. Centrally defining security policies for specific client roles, using AD Group Policy Objects and G Data ManagementServer, will provide tailored security to each client in the network. Specifically, it should be ensured that the server has appropriate fallback measures. If the server itself fails, due to hardware problems or connectivity issues, the clients should still be able to connect to network resources. In addition, it is not possible to efficiently automate backup procedures, because the server operates on its own. Server configuration and databases should be backed up to a second, separate device, or to an external storage medium.

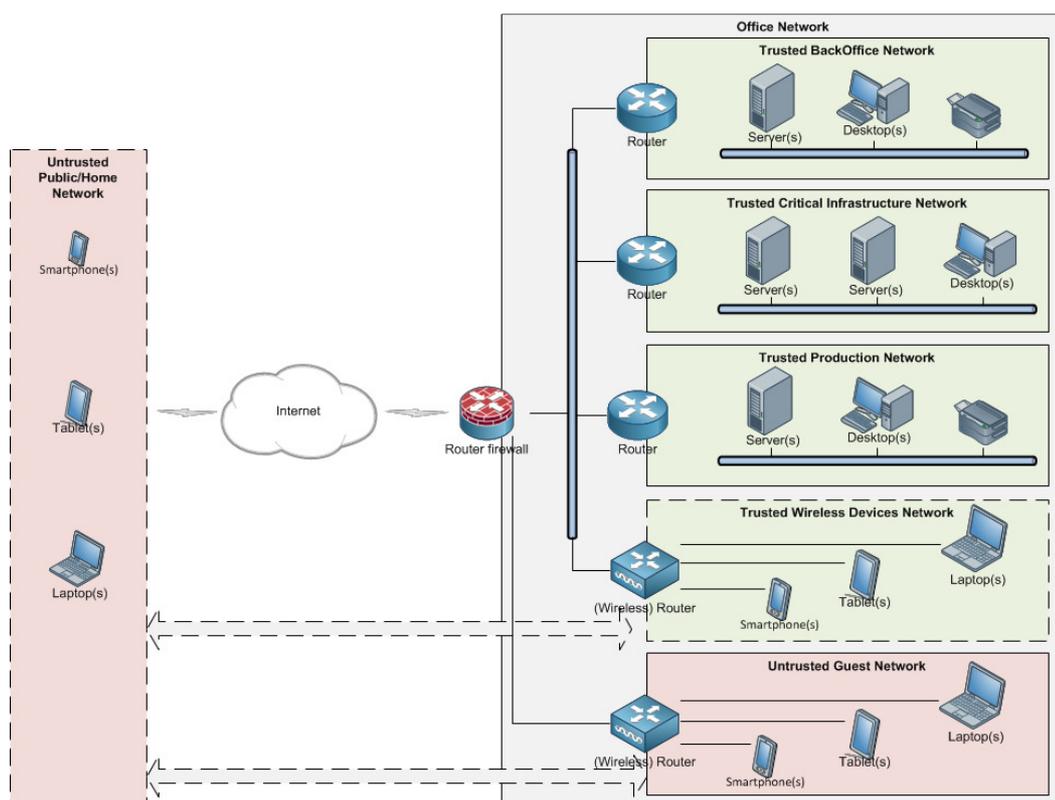


Image 2: Medium to large office network

As the number of network clients grows, a simple network will inhibit effective deployment and maintenance. A formalized network layout with its network zones and client roles will allow administrators to develop policies for the different types of clients that are served. This type of network scales easily, allowing clients numbers from ten to several hundreds. A typical medium to large office network does not rely on a single device unifying modem, router and firewall. Instead, several gateway devices are deployed to serve protected internet access. A separate network-level firewall device allows for high-performance filtering of network traffic. After the firewall, traffic is distributed to the internal network. The different network zones are physically separated by

assigning them one network device each and dividing them into different subnets. Routers and switches provide access for cabled connections to the various network zones and manage and restrict traffic between different trust zones. If they have been deployed and set up for the corporate network, mobile devices can connect to a trusted wireless devices network. For unknown devices, a separate untrusted guest network is configured.

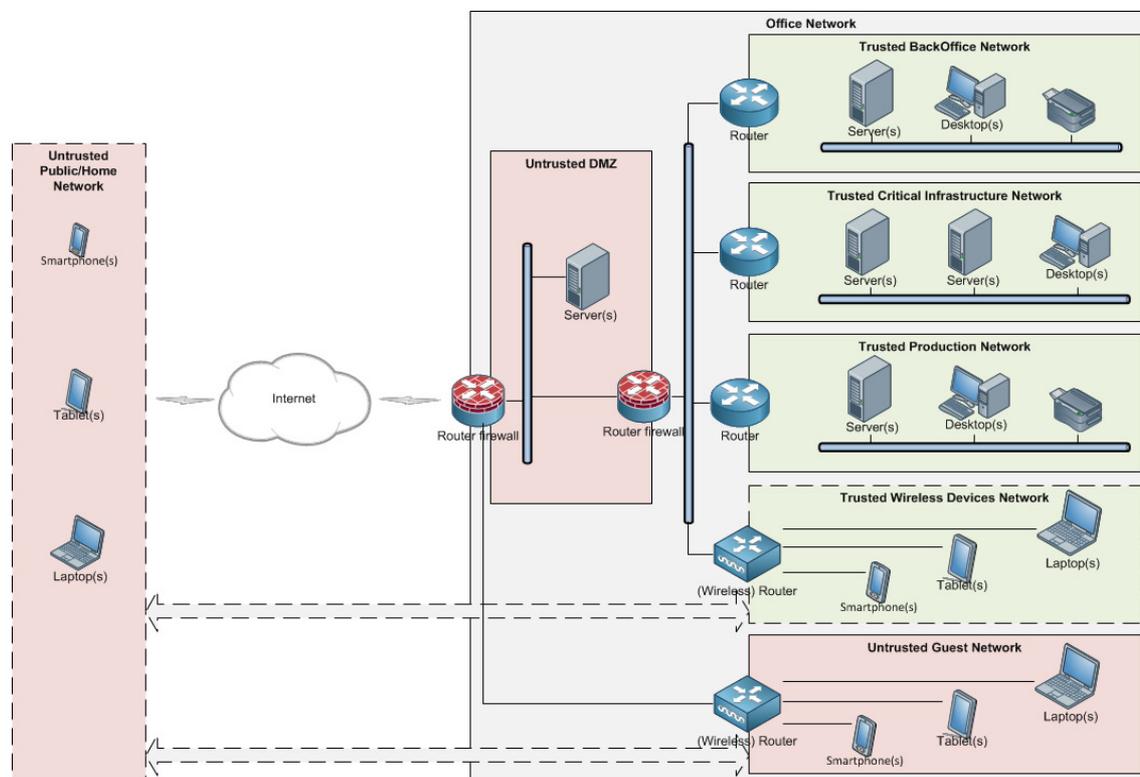


Image 3: Medium to large office network with demilitarized zone (DMZ)

It is important to distinguish between the different types of servers and the clients they are serving. In small or medium to large office networks, there will always be one or more internal servers. Whether located in a dedicated network zone or not, these server(s) will provide internal services to clients, such as file storage, print services or ERP. At the same time, one or more servers may be in use as mail servers, web servers or FTP servers, providing services to clients outside the network. Although medium to large networks can be operated with both internal- and external-facing servers in an internal zone, an extra security layer can be added when internet-facing servers are to be operated from within the network. The demilitarized zone (DMZ) is a logical sub-network which explicitly contains only those services that need to receive communication requests from outside the network. To prevent any unauthorized access to the internal network, contact between servers in the DMZ and internal network zones should be limited to those required for services to function correctly (such as e-mail). In this scenario, the two network-level firewalls play an important role, as they need to ensure that almost all outside traffic ends up in the DMZ, unless requested by an internal client. Any contact between services in the DMZ and internal services should be scrutinized carefully.

For larger businesses that have multiple office locations, network layout becomes a bit more complicated. For all locations, local services are often duplicated, while external facing servers may or may not be deployed to all locations. However, the basic premise is still the same as for a company with only one office. The internet connection will be routed through a series of network-

level hardware that filters out malicious traffic and routes it to either a server in the DMZ or one of the internal network zones.

There are many devices that do not fit the traditional networking paradigm. Employees may be using mobile devices that are sometimes used within the enterprise network, sometimes on the outside. Whether these are company-issued or not, they will need to conform with security policies when connecting to the enterprise network. The same goes for devices that are physically located outside the network, but can access resources within the network, such as clients connecting through a VPN. Security for these types of devices depends on both a hardware component, as well as software enforcement of security policies.

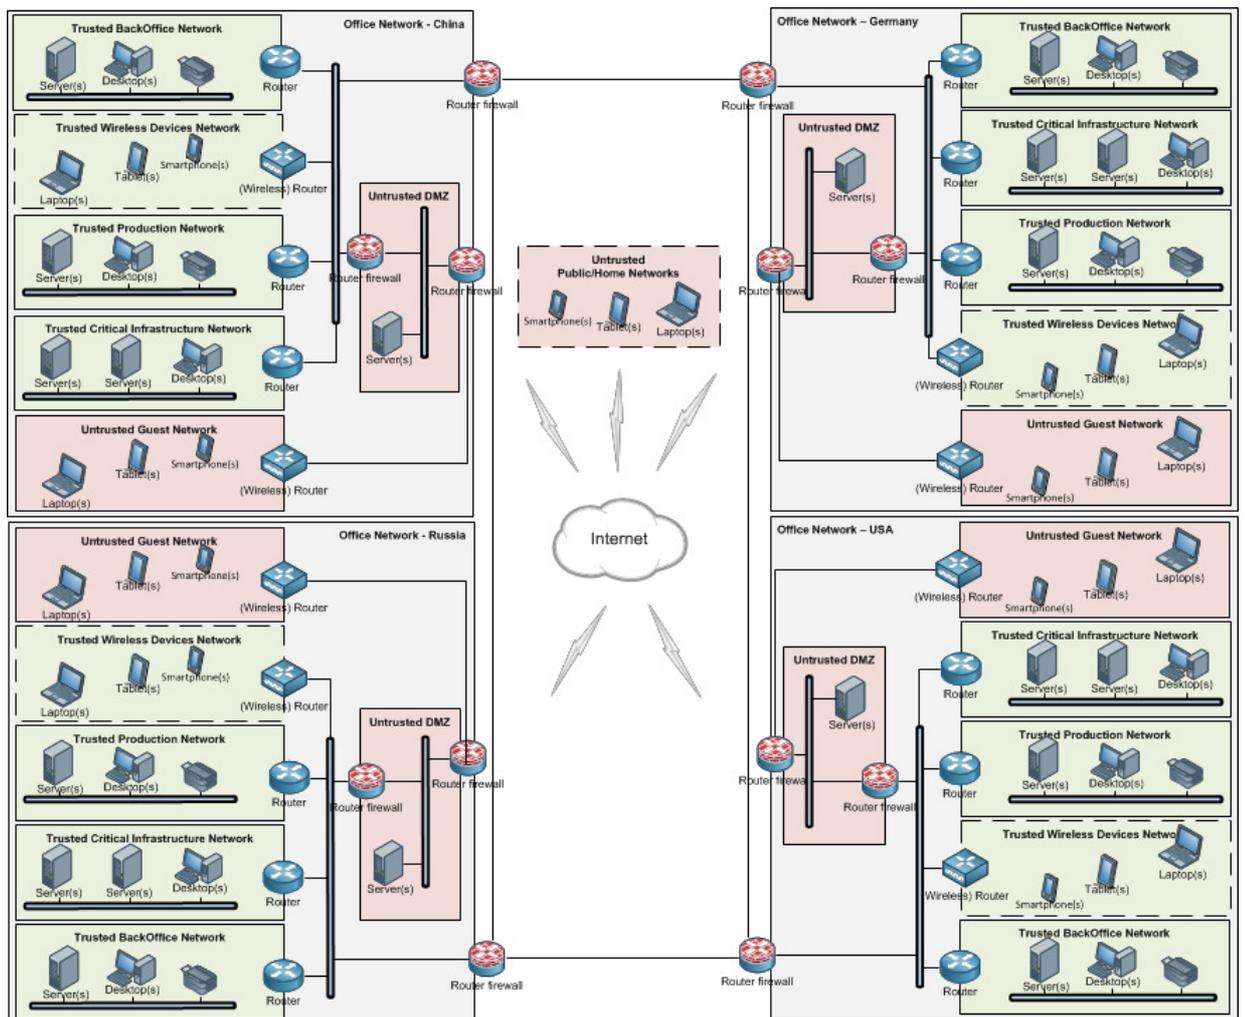


Image 4: Medium to large office network with DMZ, spread over multiple locations

With the physical network layout organized, logical units for network clients are relatively easy to set up. Ideally, each network zone will have one group of clients that all share the same role. In some cases, a network zone may have clients in more than one role, in which case multiple groups will need to be created. This client structure should be reflected in the network's organization units. When using Windows-based networks, this will usually be an Active Directory structure, where each client role is assigned its own Organizational Unit, which contains the applicable policies.

1.2. Security components

As the network diagrams in chapter 1.1 show, some security measures are implemented on the hardware level. A router with built-in firewall, or a dedicated firewall device, is a great physical security measure that filters out a lot of traffic, but it cannot provide full operational security on its own. Any network security policy must be a combination of different layers, combining solutions to build an all-encompassing security construct. Some of those layers will be positioned on the hardware level, others are software solutions. Some will provide security by analyzing network traffic, others will limit dangerous activities on clients. But none of them suffice on their own: make sure that security components cover all possible points of entry, and that they work together effectively.

The network diagram can be used as a starting point to find out which security components need to be applied and where. Other than deploying hardware-based measures, individual network devices should also be secured. For each of the devices or device categories in the network diagram, consider what type of traffic it will send and receive, which use will be made of it, and what its importance to the enterprise is. One measure is to install local security software on clients and servers, to ensure that malware cannot be spread to and from individual clients. Additionally, sensitive servers should be protected. For mail servers, inbound and outbound e-mails should be scanned for malware (with the additional benefit of adding spam filters). Chapter 3 describes the possibilities to deploy G Data software to critical network devices to provide maximum protection.

2. Choosing a solution

With network layout and client management properly organized, an informed decision can be made about which G Data business solution to install. Based on the enterprise’s needs, customers can select the solution which best applies to their network. Whether maximum security is required, flexibility, performance, or all of the above, the combinations of modules present in G Data’s business solutions will fit every network. It is important to find out which security modules are needed to optimally secure the network. For example, companies that make use of a Microsoft Exchange mail server should make sure their solution includes G Data MailSecurity and its Exchange plugin.

2.1. G Data business solutions

The chart below shows all G Data business solutions and their respective modules. Malware protection forms the baseline of G Data’s product portfolio. Every solution, starting with the entry-level AntiVirus product, contains the AntiVirus module, which combines techniques such as signature-based protection and heuristic approaches to provide excellent client-based protection. Its active hybrid technology combines two engines for optimal detection rates. The new G Data CloseGap engine optimizes performance while detecting even locally confined threats. The BankGuard module provides additional protection while using online banking services, while ReportManager allows administrators to gain insight into the state of their network and connected clients. Mobile device management integrates mobile devices into G Data’s management panel, to enable security management and privacy and anti-theft measures.

G Data ...	 AntiVirus	 Firewall	 AntiSpam	 Policy-Manager	 BankGuard	 Report-Manager	 MobileDevice Management	 MailSecurity	 Backup
AntiVirus Business	★				★	★	★		
AntiVirus Enterprise	★		★		★	★	★	★	★
ClientSecurity Business	★	★	★		★	★	★		
ClientSecurity Enterprise	★	★	★		★	★	★	★	★
EndpointProtection Business <small>(also available as Managed Service)</small>	★	★	★	★	★	★	★		
EndpointProtection Enterprise <small>(also available as Managed Service)</small>	★	★	★	★	★	★	★	★	★
PatchManagement	Add-on module for all solutions listed above								
SmallBusiness Security	★	★	★		★	★	★	★	★

Image 5: G Data software solutions

ClientSecurity adds G Data Firewall, a host-based intrusion prevention system (HIPS) which monitors client network traffic to prevent unwanted access to client systems. The client-based AntiSpam module provides protection against unwanted and infected e-mails through network-level scans as well as a plugin for Microsoft Outlook. Finally, EndpointProtection serves companies that want to

centrally manage security policies, such as media device control or internet usage time.

The additional MailSecurity and Backup modules are included with the Enterprise version of each product. MailSecurity provides spam and malware filtering capabilities for mail servers on protocol level and can be installed as a standalone gateway. Its plugin for Microsoft Exchange seamlessly integrates malware protection in the server workflow. The Backup module allows administrators to flexibly schedule file-based backups for all clients in the network to prevent data loss in case of emergency.

To save management costs on patch testing and deployment, the optional PatchManager module can be obtained. It is available as an add-on for all Business and Enterprise products and features an integrated solution to obtain and distribute the latest patches for software from all popular vendors.

For small businesses, G Data SmallBusiness Security offers a valuable set of modules. All security modules except for PatchManager and PolicyManager are included and licensed for use on up to five computers (the license can be expanded to include ten computers). Additionally, up to three non-networked PCs can be managed with the included license for G Data InternetSecurity.

To decide which solution to deploy, select the product that offers the modules that are required, keeping in mind that the PatchManager module is always available as an optional extra. Make sure that all network entities are being protected by the chosen product: (mail) servers, network clients, mobile devices. Mix and match modules to find the network's optimal solution. For example, an enterprise which requires security, but does not need to use an integrated backup solution, can choose AntiVirus Business. If the need for mail server based security arises, the MailSecurity product can be acquired separately or as part of the Enterprise edition. The most comprehensive solution is EndpointProtection Enterprise, which can optionally be complemented with the PatchManager module.

2.2. System requirements

To ensure problem-free deployment, the system requirements of G Data software should be compared to the hardware that is currently in use in the enterprise network. Make sure that all servers and clients that should be protected meet the system requirements, updating and standardizing hardware deployments if necessary. G Data optimizes its software for use on a wide range of server and client operating systems and hardware, so for most networks, deployment will be problem-free. A dedicated server is not required, but recommended for large networks.

G Data ManagementServer/G Data MailSecurity MailGateway

- Operating system: Windows 8, Windows 7, Windows Vista, Windows XP SP3 (32-bits), Windows Server 2012, Windows Server 2008 R2, Windows Server 2008, or Windows Server 2003
- RAM: 1 GB

G Data Administrator/G Data WebAdministrator/G Data MailSecurity Administrator

- Operating system: Windows 8, Windows 7, Windows Vista, Windows XP SP3 (32-bits), Windows Server 2012, Windows Server 2008 R2, Windows Server 2008, or Windows Server 2003

G Data MobileAdministrator

- Operating system: Windows 8, Windows 7, Windows Server 2012, or Windows Server 2008 R2

G Data Security Client

- Operating system: Windows 8, Windows 7, Windows Vista, Windows XP SP3 (32-bits), Windows Server 2012, Windows Server 2008 R2, Windows Server 2008, Windows Server 2003, or Linux kernel 2.6.25 or higher
- RAM: 1 GB

G Data Internet Security for Android

- Operating system: Android 2.1 or higher
- Disk space: 14 MB

G Data MailSecurity for Exchange (64-bits Exchange plugin)

- Mail server: Microsoft Exchange Server 2013, Microsoft Exchange Server 2010, or Microsoft Exchange Server 2007 SP1

G Data ManagementServer requires Microsoft .NET Framework 4.0, which will be automatically installed with ManagementServer itself. WebAdministrator and MobileAdministrator also require Microsoft .NET Framework to be installed. Additionally, as they are web services, they require Microsoft Internet Information Services (IIS) to have been installed beforehand. To log in to WebAdministrator, the local browser needs to have the Microsoft Silverlight plugin installed. For storage, G Data ManagementServer uses a SQL server. Microsoft SQL Server 2008 SP3 Express will be installed, or an existing instance of Microsoft SQL Server 2008 or 2012 can be used. When using G Data ManagementServer/G Data MailSecurity MailGateway with a local SQL database or other demanding applications on the same computer, the following recommended system requirements apply:

- RAM: 4 GB
- CPU: multicore

G Data products use the TCP/IP protocol for communication between server and clients. Certain ports need to be available on servers and clients to enable control communication and update distribution. Make sure to configure network-level monitoring software and firewalls to allow traffic on those ports. In case of port conflicts, a manual reconfiguration of some of the port numbers is possible (see chapter 17.1.2).

Main server/secondary server

- Port 7161 (TCP): Communication with clients and subnet servers
- Port 7182 (TCP): Communication with G Data Administrator
- Port 7183 (TCP): Communication with mobile clients
- Port 7184 (TCP): Communication with mobile clients (distribution of mobile client installation files)

Subnet servers

- Port 7161 (TCP): Communication with clients and (subnet) servers

Clients

- Port 7167 (TCP): Communication with (subnet) servers
- Port 7169 (TCP): Communication with clients (peer-to-peer update distribution)

MailSecurity MailGateway server

- Port 7182 (TCP): Communication with G Data MailSecurity Administrator

MailSecurity Exchange plugin

- Port 7171 (TCP): Communication with G Data Administrator

2.3. Licensing

Once a solution has been decided upon, information about licensing can be obtained from the G Data website¹ or an official sales partner. In general, business licenses are available for deployments of 5 clients and more. The price of individual licenses depends on the solution chosen, eventual add-on modules, and the number of clients being deployed.

¹ See www.gdatasoftware.com.

3. Installation scenarios

After deciding on an appropriate G Data solution for the network, deployment needs to be planned. G Data products make use of the client-server model. A central server application manages any number of clients in the network, optionally supported by a secondary server and one or more subnet servers. On each client machine, client software manages security, backup, patching and other processes. Deployment first focuses on setting up one or more servers, which are then used to deploy client software to network machines. Whether to use one or more servers depends on the network layout. Small networks can be managed by one local ManagementServer, while deployments on large networks or at different locations can take advantage of a setup with multiple ManagementServers, remotely managed by a central MasterAdmin installation.

G Data solutions can be deployed as a locally managed product, or as a ManagedService. The former gives administrators the flexibility they need to configure the product to their needs at any time, but requires time and effort in order to get acquainted with the product and tailor it to the network's needs. The latter offloads management duties to a ManagedService partner, who remotely configures and administers the solution, which does not require any interaction from the local administrator.

Configuration of G Data solutions takes place using G Data Administrator, G Data WebAdministrator and/or G Data MobileAdministrator. This range of tools offers configuration possibilities including local administration, remote administration, browser-based access and mobile configuration. For more information about the administration tools and their use cases, see chapter 5).

3.1. Local deployment

The client-server model of G Data can be applied to every network configuration. An installation of one or more server components (G Data ManagementServer and its secondary or subnet server(s), and G Data MailSecurity) is combined with a client software component on every client (G Data Security Client and G Data Internet Security for Android). The various component types and the deployment possibilities for the central ManagementServer component will be discussed, based on the network layout diagrams presented in chapter 1.1.

3.1.1. Network components

G Data solutions consist of multiple network components. Depending on the network layout and the network requirements, various components may be installed. ManagementServer is central to the client-server based protection concept and will take a central place in every deployment. There are various ways to deploy it; chapter 3.1.2 provides examples for networks ranging from small offices to large corporate scenarios.

The client component (G Data Security Client) will be installed on network PCs. Security Client will provide several layers of protection: its firewall will check inbound and outbound network traffic, while its file system monitor module checks all files for malware. For clients that are not in the enterprise network, SmallBusiness Security includes three licenses for G Data InternetSecurity, a standalone security solution. Mobile devices can be protected by G Data Internet Security for Android.

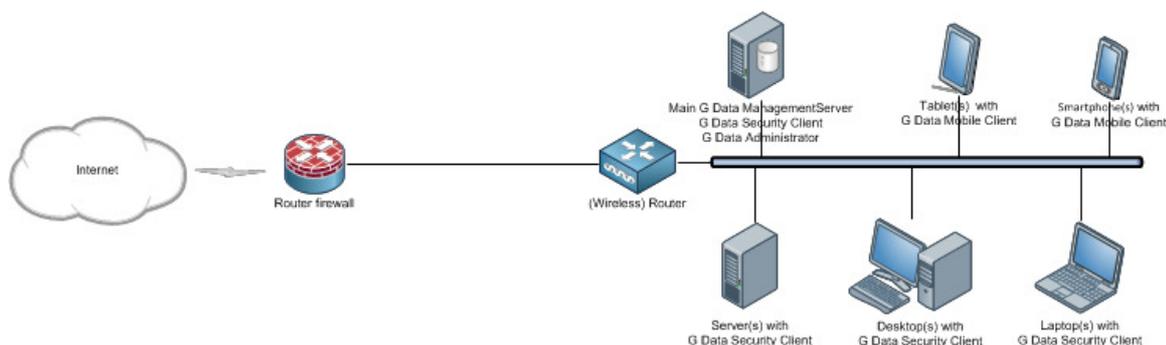


Image 6: Small office network deployment

For networks that host their own mail server, the Enterprise solutions and SmallBusiness Security include MailSecurity, the mail server protection component, which scans mails for malware and spam. It can be installed as a standalone product on its own server, filtering traffic before it hits the mail server, or as a plugin for the popular mail server Microsoft Exchange.

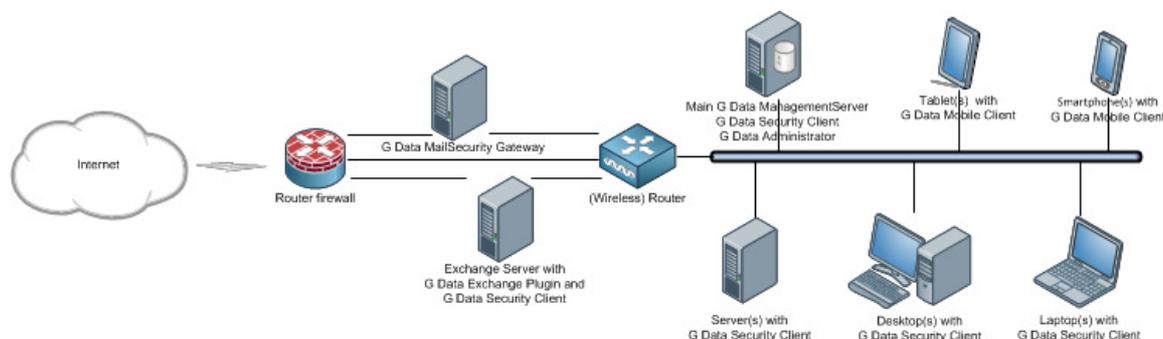


Image 7: Small office network with mail server deployment

3.1.2. ManagementServer deployment

Which deployment type should be chosen for ManagementServer depends entirely on the network. Current and future properties of the network should be taken into account, such as the type of infrastructure, the number of clients, and the types of client access. The scenarios below provide an indication of what a G Data deployment looks like for several types of networks. Because of the modular nature of the solutions, deployment can be tailored to each network's circumstances. In addition, deciding on one deployment scenario does not mean that the installation cannot be adapted at a later time. If the network grows beyond a certain number of clients, components can be moved around the network and additional (subnet) servers can be deployed.

G Data solutions' main installation will be the server component, G Data ManagementServer. For performance reasons, it is not recommended to install ManagementServer on one of the clients that is also used for daily tasks, but it is possible: ManagementServer does not require a server operating system, so it can easily be installed on Windows XP, Windows 7 or any other supported Windows desktop operating system – a viable alternative for small office networks. When using a dedicated server, it does not need to have a server operating system or dedicated server hardware, but the machine should not be used for any other purpose than to host ManagementServer. Be very careful when installing other services on the same server, especially database, mail security, domain controller or web server services. The more services are run from the same machine, the more all of them will be impacted, leading to delays during peak traffic hours. Similarly, the number of clients being served has direct effects on server performance. For networks with a small number of clients,

the load will be relatively small, allowing for more services to be run from the same server. Larger networks may hit the boundaries of server performance sooner, requiring services to be moved to a dedicated server or split over multiple servers.

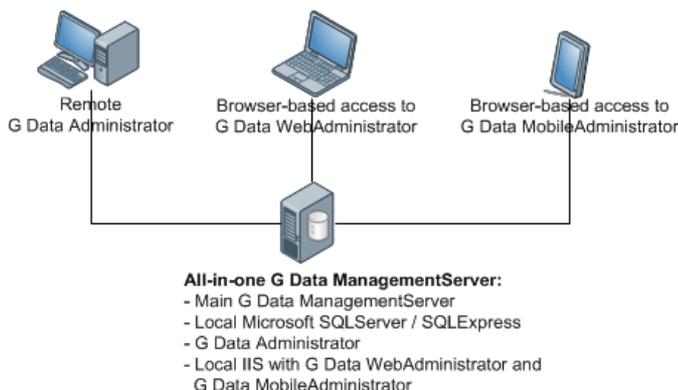


Image 8: All-in-one ManagementServer deployment

On the server level, the easiest type of deployment is an all-in-one deployment. For networks that have a relatively small number of clients, or that do not have a dedicated server, all components of the ManagementServer deployment can be installed on the same machine. This includes the ManagementServer installation, as well as a local installation of Microsoft SQL Server or the included Microsoft SQL Server 2008 SP3 Express, and a local installation of G Data Administrator. When remote administration capabilities are required, Microsoft Internet Information Services (IIS) and G Data WebAdministrator or MobileAdministrator can be deployed to the same machine as well.

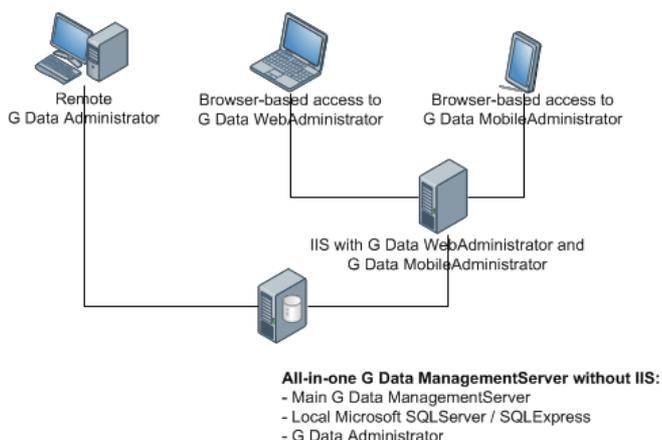


Image 9: ManagementServer deployment with dedicated web server

Installing ManagementServer and its database and administration components on one machine results in a server that is easily manageable, but prone to performance issues if the network grows. Optionally, IIS and administration components can be deployed to their own web server. While they are not very performance-intensive, a network that already has its own IIS server does not need an additional IIS installation on the ManagementServer. WebAdministrator and MobileAdministrator can be easily deployed to any existing IIS server, starting at IIS 5.1 (Windows XP). As detailed in the chapter on network layouts (see chapter 1.1), the web server can be located either in the network itself, or in the DMZ.

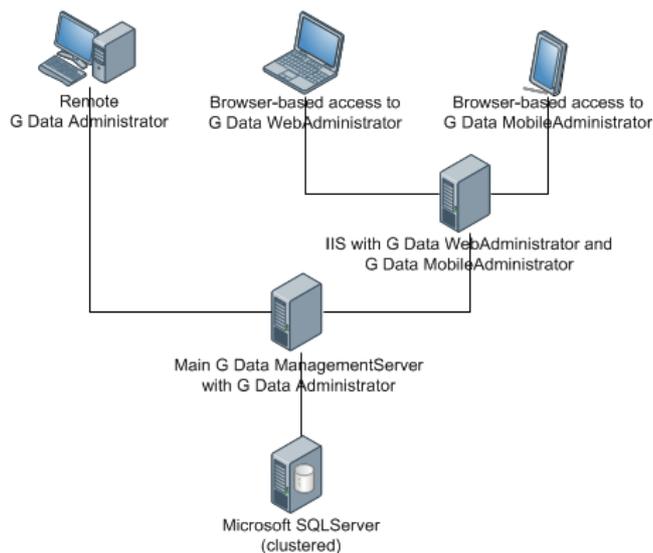


Image 10: ManagementServer deployment with dedicated (clustered) SQL server

Networks with more than 1000 clients will quickly make an all-in-one ManagementServer deployment reach its performance limits. Although ManagementServer has built-in load limit capabilities (see chapter 7.5), altering the deployment can be a far more effective solution. A large part of server load can be taken care of by deploying a (clustered) Microsoft SQL Server. Hosting the ManagementServer database on a dedicated SQL server improves the performance on the ManagementServer by offloading its database transactions. Although it is possible to use a dedicated SQL server right from the beginning, the most common use case will be a migration scenario, when it turns out the number of clients has grown too large to be managed by a database on the ManagementServer itself. That process is simple: after installing and configuring Microsoft SQL Server on its dedicated server, the GdmsmsConfig.exe tool can be used to migrate the database. See chapter 17.2.1 for more information about database migration.

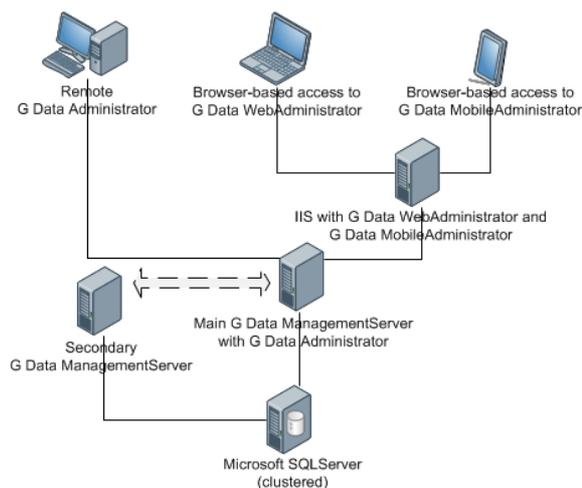


Image 11: ManagementServer deployment with secondary ManagementServer

A medium to large office network can profit from installing a secondary ManagementServer as a fallback measure. The secondary ManagementServer is installed on another server and functions parallel to the primary ManagementServer. If the main ManagementServer is unavailable for more than an hour, clients will connect to the secondary server to obtain updates. The main and secondary ManagementServer share the same database, so this option can only be used if the

database is hosted on an external Microsoft SQL Server. While the two servers do share the same database, they obtain updates from G Data's update server autonomously. This provides extra redundancy in case one of the servers loses internet connectivity. In combination with a dedicated, clustered SQL server, the secondary server deployment type provides a high reliability, preventing problems when hardware issues occur.

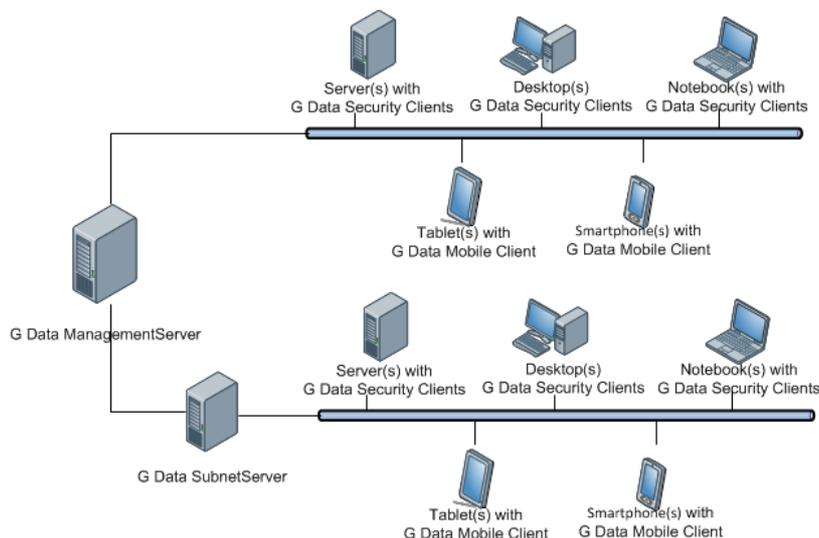


Image 12: ManagementServer deployment with subnet server

To ease the load, G Data ManagementServer supports the deployment of one or more subnet servers that manage subsets of clients, reducing the amount of performance required from the main ManagementServer. This is especially useful for large networks and networks spanning several branch offices. A subnet server is an installation of ManagementServer that takes care of a subset of clients. It reduces network load, as the ManagementServer only has to contact the subnet server, which will then serve its clients automatically. Using subnet servers, a single ManagementServer deployment can service thousands of clients effortlessly. Subnet servers are typically installed after deploying the main ManagementServer and its clients – any client can be configured as a subnet server through local or remote installation. The main ManagementServer does not need to be located in the same physical network as the subnet server. Branch offices can be managed by a subnet server, which connects to a central ManagementServer. The same setup can be used by service partners that manage a customer deployment. It does incur some limitations: direct communication between a ManagementServer and its clients is limited when it takes place over the internet. For example, the client folder selection used in some file selection dialogs does not work.

In some situations, multiple ManagementServers need to be managed. A large company can divide its clients over multiple ManagementServers; service partners may manage ManagementServers for multiple customers. In this case, administration can be centralized by adding all ManagementServers to a MasterAdmin deployment. The ManagementServers function independently from each other, but are centrally managed by a MasterAdmin installation. This requires the appropriate port forwarding settings to make sure that the servers can be reached over the internet. Otherwise, it looks much like a regularly managed deployment. Using MasterAdmin, administrators have access to all modules and settings. More information about administering servers using MasterAdmin can be found in chapter 5.4.

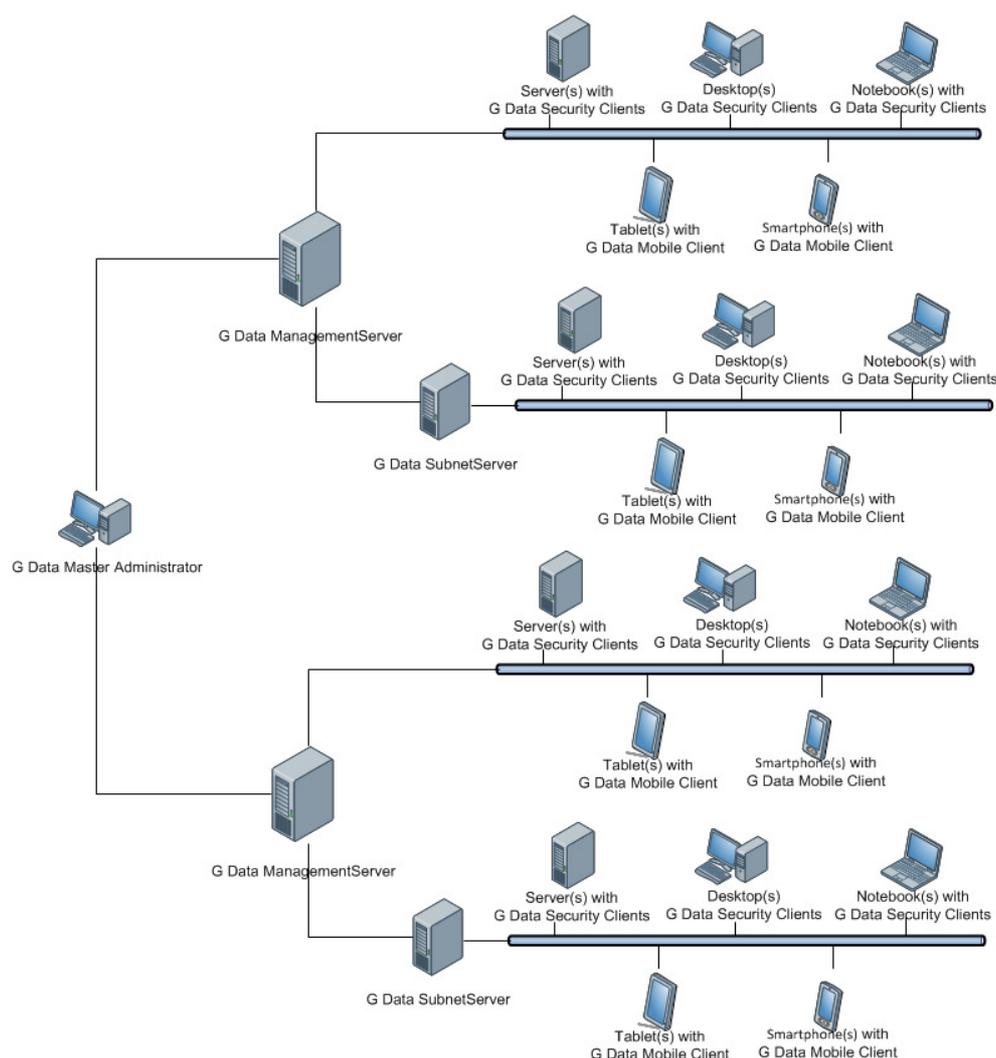


Image 13: ManagementServer deployment with MasterAdmin

Regardless of the size of the network, the effort that has been put into organizing the network will be rewarded when using a G Data solution. Its security modules have been organized around the concept of client groups. Security settings, scan or backup tasks, and every other aspect of security can be applied to single clients or to groups of clients. Groups can be manually created and contain any number of clients, either by mirroring a network zone and/or client role, or by grouping them according to other attributes. For networks that have organized clients in Active Directory Organizational Units (OU), the effort will be reduced even further, as client groups can be linked to OUs to automatically inherit their list of clients.

3.2. ManagedService

An alternative to local deployment and administration is the use of G Data's ManagedService product. Administrators can choose to obtain G Data EndpointProtection as a ManagedService, with a minimum of 75 clients. G Data partners that offer this service help their customers by taking all work out of their hands. Deployment as well as management of the complete solution will be taken care of by the partner. For customers, the ManagedService approach saves a lot of time and effort during both the deployment phase and the management cycle. Employees do not need to be trained to manage the security solution, and administrators can focus on other management tasks.

For partners, ManagedService helps reach those customers who would not have considered using a locally managed G Data solution.

Management tasks can be carried out remotely, without any effort by the local administrator. The ManagedService proposition not only greatly simplifies deployment, but it also represents a transparent license model. Partners charge enterprises monthly for the exact number of clients in use, allowing monthly fluctuations as clients are added or removed.

On the client side, a ManagedService deployment will look and function just like a local deployment as detailed in the rest of this guide. However, no effort needs to be put in management. All tasks will be taken care of remotely by the partner that provides the service. Partners administer ManagedService networks remotely using MasterAdmin (see chapter 5.4).

4. Deployment

Whether a deployment is carried out by a local administrator or remotely by a ManagedService partner, several scenarios can be followed, depending on whether the solution will be installed as an upgrade or as a clean installation. In any case it is recommended to run a trial deployment before affecting the physical network. A virtual network or a subset of the physical network can be used to install the G Data solution on server and clients, to see if any problems arise. Make sure that the trial is run across a group of clients that represents all client roles (e.g. IT, back office, R&D) to accurately test deployment results across the whole network. The test run as well as the actual deployment of a G Data security solution should be carried out in accordance with existing corporate deployment policies, if available. For small organizations, this can mean a fairly straight-forward process, but larger enterprises may have to develop a project plan, documenting deployment planning, risk assessment and more.

Every deployment is carried out in several stages. Because G Data products have been designed as server-client solutions, the first step is installing and setting up the server component(s). After the installation, server settings (such as update schedule and distribution) and default client settings should be configured to make sure that clients are properly protected upon rollout. After completing the actual deployment phase, check if all clients have been properly deployed. With all clients regularly connecting to the main ManagementServer, settings can be customized for each client.

G Data MailSecurity, acquired as optional module or as part of an Enterprise solution, is a stand-alone product. As such, it has a separate installer, but furthermore, the configuration is slightly different from the security products. During the installation, MailSecurity will be configured according to its place in the network (see chapter 3.1.1). Afterwards, its protection and anti-spam measures can be customized completely.

4.1. Preparation

Before deploying any server or client installations, it is essential that the systems meet a certain baseline. Run Windows Update on the server(s) where G Data ManagementServer will be installed and install all available security updates. If any frameworks are present, such as Microsoft .NET Framework, these should be updated as well. Finally, clients that are to be installed remotely need to be prepared with the proper file and folder access permissions (see chapter 4.7.1).

Before installing a security solution to any machine, it should be ensured that there is no malware present. Existing malware can interfere with the installation of security software, compromising system security. For this purpose, the installation medium of every G Data business solution doubles as a bootable medium with its own Linux-based operating system². When starting from the boot medium, the G Data AntiVirus boot environment allows administrators to carry out a complete malware scan of all local hard disks. This makes sure that there are no traces of malware left on the system. For smaller networks, it is recommended to use the boot medium on all servers and clients before installing G Data software. For larger networks, however, this is not always an option: each client has to be booted from the boot medium manually, and a full system scan will take some time, leading to a long downtime for servers. In that case, the scan could be executed on only a limited number of systems, such as high-priority or high-risk clients and on servers that are suspected to

² If the original physical installation medium is not available, the digital (ISO) version can be used to run the G Data BootCD Wizard, which creates a bootable G Data BootScan medium (CD, DVD or USB stick).

have been infected with malware.

4.2. Clean installation

When no G Data solution has been previously deployed to the network, all components can be deployed as a clean installation. This means that the main server, potential secondary or subnet servers, local database instance, and clients will all be set up from scratch. Using the default settings, this is an easy procedure. The server component(s) will be installed first, followed by administration tools, clients and mail server security. With minimal configuration effort, a clean installation can be carried out within a few hours. However, as with any deployment, installing a G Data solution requires that the administrator knows how the network is structured. Components should be installed on the appropriate targets and a basic configuration should be carried out. Chapters 1, 2 and 3 provide guidance on network structure.

4.2.1. Server

During the ManagementServer installation, several options need to be set, such as the server mode (main ManagementServer, secondary ManagementServer or subnet server), and the database configuration. These decisions fully depend on the network structure and the resulting deployment decisions.

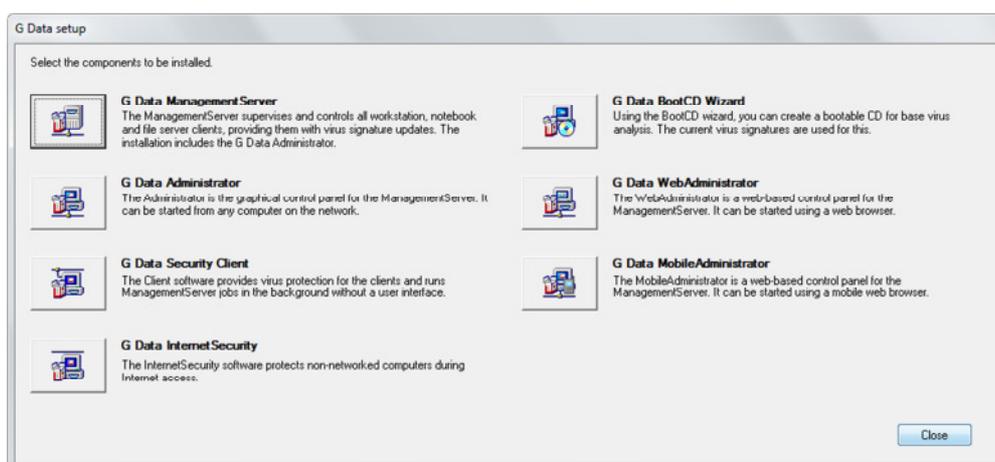


Image 14: G Data installation DVD, product selection

The main ManagementServer installation is the first component to be installed. It is the central server that will coordinate all clients, potentially supported by a secondary server or one or more subnet servers. ManagementServer can be installed from the G Data installation medium by selecting G Data ManagementServer in the component selection window. While the installation wizard itself is very straightforward, two steps need additional attention. ManagementServer can be installed as a main server, a secondary server or a subnet server. The first installation will be a main server. If the network layout requires the deployment of a secondary server or subnet server(s), these can be installed on their respective machines after finishing the installation and initial configuration of the main server. When installing a subnet server, the computer name of the main server needs to be entered, to allow the subnet server to contact the correct main server. Alternatively, subnet servers can be installed remotely through G Data Administrator.

After selecting the server type, the database type has to be selected. For networks up to 1000 clients, a local installation of Microsoft SQL Server 2008 SP3 Express will suffice. If the Express option

is selected, the installation wizard will automatically install and configure the local database server instance and required database. For networks with more than 1000 clients, or scenarios where a secondary server will be used, the database should be configured to use an existing Microsoft SQL Server instance on a separate, dedicated server.

After the installation procedure is completed, the installation wizard will ask for information about product activation. When the product has not been activated before, the license key can be entered to automatically activate the software and request a user name and password. User name and password are then saved in the product's configuration, so that updates can be downloaded automatically. If the product has been activated before and the license is still valid, the user name and password can be entered manually. Alternatively, activation can be postponed. Manual activation afterwards is supported (see chapter 4.5), but not recommended. Without an activated product, only the very basic functionality will be available. Even if the product has been bought as an EndpointProtection or ClientSecurity edition, only the functionality of the AntiVirus edition will be available until the product is activated. Moreover, without activation no program file and virus signature updates can be downloaded. This severely limits the effectiveness of the various protection layers.

The ManagementServer installation will have installed Microsoft .NET Framework 4.0 if it was not already present on the server. It is recommended to use Windows Update after installing ManagementServer to check for updates for Microsoft .NET Framework 4.0. A reboot after installation is recommended in any case.

Once the installation for the main server has been carried out, the next step depends on the deployment wishes. First, the main server should be configured (see chapter 4.4). If a secondary server should be deployed to the network, the ManagementServer installation wizard should be run on the respective server to install it. If subnet servers are to be deployed, this should take place after deploying the clients; subnet servers can easily be remotely deployed through G Data Administrator. For all servers, it should be made sure that the TCP ports that are used for communication are available. If a network-level or software firewall is in use, a number of ports should be opened. See the system requirements (chapter 2.2) for more information about the TCP ports that G Data server and clients use, and how to change them if there are any conflicts.

4.2.2. Administration

After the servers are up and running, consider the administration requirements. Every ManagementServer installation includes a local installation of the G Data Administrator tool, which allows administrators to configure ManagementServer. If the server is going to be configured remotely, which is strongly recommended, there are several possibilities: installing G Data Administrator on another client, installing G Data WebAdministrator to configure ManagementServer remotely from any browser, or installing G Data MobileAdministrator to administer the server using a smartphone or tablet. The options for remote administration are documented in chapter 5.

4.2.3. Clients

As soon as ManagementServer and Administrator have been deployed, G Data Security Client should be distributed to all network clients. To protect the server itself from malware, it is recommended to install G Data Security Client on the server as well. Like other network clients, the

remote installation possibility is the easiest way to do this. See chapter 4.7 for more information about client deployment.

4.2.4. Mail server security

If the deployment includes G Data MailSecurity (part of SmallBusiness Security and G Data Enterprise solutions), it should be installed after ManagementServer has been deployed. Using its dedicated installation medium, MailSecurity can be installed as an Exchange plugin, if Microsoft Exchange is being used as mail server. Alternatively, MailSecurity can be installed as a standalone gateway product. MailSecurity MailGateway functions as an extra security layer in front of the actual mail server, processing all inbound and outbound mail. This can be achieved in two ways: by installing MailGateway on the mail server itself, or by setting it up as a gateway on a different server.

4.2.4.1. Exchange plugin

The Exchange plugin of G Data MailSecurity can be installed on any Exchange Server installation from Exchange Server 2007 SP1 onwards. It reports to G Data ManagementServer. The installation wizard will either configure MailSecurity to use an existing ManagementServer or install a new, local ManagementServer – because of system performance, using an existing ManagementServer is recommended. After the installation concludes, the Exchange plugin will register itself with the ManagementServer. By logging in to the ManagementServer using G Data Administrator, all settings of the Exchange plugin can be managed.

4.2.4.2. Standalone installation

If another mail server is in use within the network, or if Exchange should be protected by a dedicated gateway including malware protection and spam filter, MailSecurity MailGateway can be installed as a standalone product. If the existing mail server can deal with the CPU and RAM load, installing MailGateway on the same machine is possible. This has the advantage that no changes to the mail server's IP address need to be made. However, the mail server software then needs to be reconfigured to use different ports for inbound and outbound mail. Alternatively, MailGateway can be installed on a dedicated gateway server, filtering e-mail before it reaches the mail server. See chapter 16.2.1 for more information about deployment types and port settings for MailGateway.

The installation wizard of MailGateway is straightforward. Optionally, it can install a local database server (Microsoft SQL Server 2008 SP3 Express). This enables statistical assessment of e-mail messages (see chapter 16.2.1) and greylisting (see chapter 16.2.5.4), but is not required.

Regardless of the server on which MailGateway is installed, G Data Security Client should be deployed to the same server first. Not only will this protect the local server's file system from malware, but MailGateway will also automatically integrate Security Client's virus signatures into its malware scan.

MailGateway's settings are configured using G Data MailSecurity Administrator, which is installed automatically with MailGateway. Like G Data Administrator, which can configure ManagementServer remotely, MailSecurity Administrator does not need to be installed on the same server as the MailGateway component itself. Using the installation medium, MailSecurity Administrator can be installed on any network client that can access the MailGateway server. To allow MailSecurity Administrator to contact the MailGateway server, access on TCP port 7082 should

be allowed.

4.3. Upgrade installation

For existing deployments of G Data ManagementServer and its clients, the upgrade path is usually very simple. There are two ways of updating a main ManagementServer installation. The most straightforward way is to use the Internet Update tool to download and install the new version. This is an automatic process that only requires the ManagementServer service to be restarted afterwards. If a secondary server has been deployed, it will be automatically informed by the main server as soon as the main server update has been deployed and subsequently update itself. Subnet servers will also be updated automatically. After the servers have been updated, the clients will be served with updated software. This happens as part of the regular software update mechanism: automatic program file updates, or manual distribution (see chapter 7.3.2). For larger networks, peer to peer update distribution is recommended to make sure there is no performance loss on the server due to the large amount of traffic. As with all software distributions, it is important to make sure that there are no compatibility problems. Using staged software distribution, the updated client can be distributed to a small group at first, before generally deploying it across the network. Alternatively, a manual distribution across a small, representative test group can give insight into possible problems.

For version upgrades where a direct update is not possible, the main ManagementServer should be uninstalled before the new version can be deployed. The existing database should not be removed; it can be used in the new version and will be converted if required. During the installation of the new version, the database will be automatically recognized and at the first start, G Data Administrator will show all clients and settings as they were before. The same process applies to a secondary server. Subnet servers will be updated automatically after the main ManagementServer has been updated. For an extra layer of security, the old database should be backed up before removing the ManagementServer (see chapter 4.6). In case of problems during the installation the new version, a downgrade or reinstall can be carried out before restoring the original database.

A gateway installation of MailSecurity can be updated through its MailSecurity Administrator interface. Under UPDATE, the current installed version of MailSecurity and Administrator are shown. Click Software update to initiate the update procedure, which is carried out seamlessly, similar to the way ManagementServer is updated by its Internet Update tool. The Microsoft Exchange plugin of MailSecurity updates itself as soon as a new version becomes available, but only if its governing ManagementServer has already been updated to the new version as well and the option UPDATE PROGRAM FILES AUTOMATICALLY has been enabled. Alternatively, G Data Administrator can be used to manually initiate an update.

4.4. Initial configuration

After finishing the installation of G Data ManagementServer and the other components, the initial configuration of server and client settings will take place. Settings can be configured using the G Data Administrator tool, which has been automatically installed on the same machine as G Data ManagementServer. If the server is to be configured remotely, set up G Data WebAdministrator first or install G Data Administrator on a network client with access to the main MMS (see chapter 5). Regardless of the application that is used, administrators can log in with a local or domain administrator account.

There are a few settings that need to be configured before deploying the clients. The `SERVER SETUP WIZARD`, which is launched when first logging in to G Data Administrator, helps take care of the most essential settings and can be run in G Data Administrator as well as G Data WebAdministrator. After the initial setup, the wizard can still be launched from the `ADMIN` menu. Additionally, most options are available separately through the various configuration menus of G Data Administrator.

The first step of the Server setup wizard covers client deployment. For the initial setup, this step can be skipped (in chapter 4.7, the various methods of client deployment will be covered). Important is the configuration of automatic internet updates. Note that these settings concern the download of virus signature updates and program file updates from G Data's update server to the MMS. The subsequent distribution of updates to network clients can be configured at a later stage. If the product has been registered during setup, user name and password will already have been filled in. Otherwise, the tool Internet Update can be used to register the product and obtain a user name and password (see chapter 4.5). The update schedule for client virus signatures and client program files needs to be configured to meet the network's needs. Check the two schedules and alter if required. For servers that make use of a permanent internet connection, hourly update checks are recommended. For hourly updates, the exact update time can be defined by entering minutes past the hour. For example, program file updates can be checked at 15 minutes past the hour, while signature files will be checked at 45 minutes past the hour. Avoid scheduling both update checks at the same time, because this can cause a temporary performance drop.

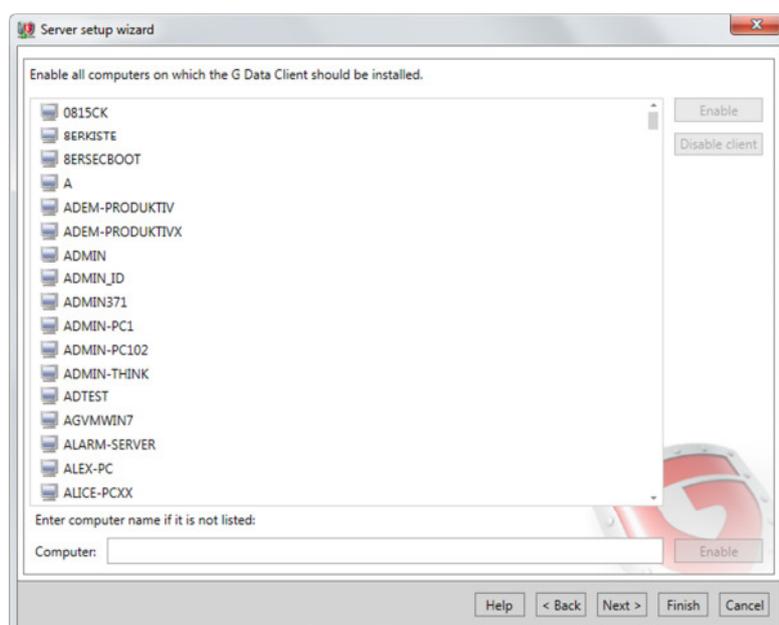


Image 15: G Data Administrator, Server setup wizard, Enable clients

The next step configures e-mail reports. They do not need to be configured in order to conclude the `SERVER SETUP WIZARD`, but doing so is recommended. See chapter 6.2 for more information about configuring e-mail reports. Even when not configuring e-mail reports right away, this is the right time to enter e-mail server settings and define recipient groups. The cogs icon leads to the `E-MAIL SETTINGS` window, where an SMTP server can be defined. Enter a valid SMTP server and port (usually 25) and a sender e-mail address. This e-mail address will also be used as a reply-to address when submitting items to G Data Security Labs. Under `MAIL GROUPS`, groups of recipients can be defined that will be later used for e-mail reports and other functions. Groups like Administrators, Management or Technical make sense: every entity that should be kept in the loop with notifications of significant

server events or emergency notifications. Mail server and recipient group settings can be edited afterwards through the `SERVER SETTINGS` window.

The last step of the `SERVER SETUP WIZARD` features some basic settings for mobile device management. Enter a password with which mobile devices will have to authenticate with the `ManagementServer`. To be able to use emergency actions, you have to enter the `SENDER ID` and `API KEY` of your Google Cloud Messaging account. Free accounts for this push notification framework can be registered at code.google.com/apis/console. See chapter 11.6 for more information about configuring a Cloud Messaging account. The settings can be edited later by opening the `MOBILE` tab of the `SERVER SETTINGS` window.

After setting up the configuration for mobile devices, the `SERVER SETUP WIZARD` concludes by initiating client installations (which will be discussed in chapter 4.7). However, there are some additional settings to be configured before deploying client software to the network. Under the `OPTIONS` menu, the `SERVER SETTINGS` window offers vital options regarding client synchronization and updates. On the `SYNCHRONIZATION` tab, configure the `MAIN SERVER SYNCHRONIZATION INTERVAL FOR CLIENTS`. This value determines how often clients check back with the MMS to see if new updates or settings are available. The more clients that will be deployed, the higher the network load will be when planning regular synchronizations. An acceptable value is 5 minutes, which can be reduced if the network only features a small number of clients, and increased if the number of clients causes load spikes. The `SOFTWARE UPDATES` tab is used to enable staged distribution for client software updates. For networks with more than ten clients, staged distribution reduces the server and network load spike that can occur when client program file updates are available. In addition, it will let clients be grouped so that critical systems can be updated later, with the first stages acting as testing platforms. To avoid high server load in case of signature updates, peer-to-peer distribution can be enabled on the `VIRUS DATABASE` tab of the `INTERNET UPDATE` menu. This allows clients to serve as update servers to each other, distributing program file and signature updates. The peer-to-peer update distribution is managed completely by the main MMS and will function without interaction, provided the appropriate client port has been configured (see chapter 2.2). Advanced settings can be altered by editing the appropriate configuration file (see chapter 16). See chapter 7.3.2 for more information about staged distribution and peer-to-peer distribution.

4.5. Server updates and registration

The `SERVER SETUP WIZARD` configures `ManagementServer` to regularly check for updates for client program files and virus signatures. However, updates for the `ManagementServer` itself cannot be automatically installed. Since server updates require a restart of the `ManagementServer` background service, they will always have to be carried out manually. If a program update for `ManagementServer` is available, customers will usually be notified by e-mail. Alternatively, the `Internet Update` tool can be used to check for updates manually.

To check for `ManagementServer` updates, click `Start > (All) Programs > G Data > G Data ManagementServer > Internet update`. The tool offers several functions that are all related to updating `G Data ManagementServer`. To check for updated program files, click `UPDATE PROGRAM FILES (SERVER)`. This will check the `G Data` update servers for updated `ManagementServer` program files, and initiate the installation procedure if any updates are available.

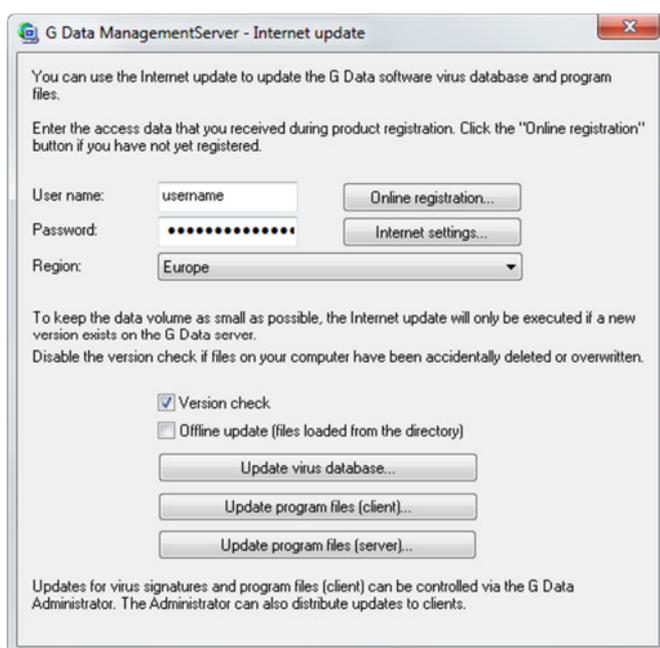


Image 16: Internet Update

In order to perform any updates, whether through Internet Update or through the automated processes of G Data ManagementServer, user name and password need to be entered. Those data can be obtained by completing online registration, which is usually carried out during the installation of ManagementServer. The wizard will then automatically request a user name and password and save them. However, if the software is not registered at that time, online registration can be manually completed afterwards as well using the Internet Update tool.

Click ONLINE REGISTRATION to open the registration form. Fill out the form, and be sure to correctly enter the registration number. Clicking LOGIN will submit the data to G Data and generate a user name and password. Be sure to make note of the user name and password somewhere, as the registration number can only be used once. When ManagementServer is reinstalled, user name and password can be entered in the installation wizard to enable updates.

In addition to ManagementServer updates, Internet Update can also download client program file updates and updated virus signatures. This functionality is identical to that of G Data Administrator, with one addition: updates can be loaded from a local folder, in case a server without internet access should be updated (see chapter 7.3.1).

4.6. Server database backup and restore

As with all types of data, it is recommended that G Data ManagementServer's database should be backed up regularly. In case of a hardware failure or other data storage issue, a recent backup should be available to quickly get ManagementServer up and running. Depending on the choice made during installation, the database is either stored locally as a SQL Server Express instance, or on a remote SQL Server installation. Using the GdmmsConfig.exe tool, a complete backup of the database can be made, regardless of where it's stored.

GdmmsConfig.exe is located in the installation folder of G Data ManagementServer, typically C:\Program Files (x86)\G DATA\G DATA\AntiVirus ManagementServer. Its interface displays several parameters related to the SQL server deployment. Server name, authentication and database data

are automatically read from ManagementServer's configuration. The server name can be a local SQL Server Express instance, or a remote SQL Server. This requires care with defining the login data. The user account needs to have access permissions for the database server, as well as write permissions for the database backup folder. The database name is GDATA_AntiVirus_ManagementServer by default, but can be customized for advanced migration or restore scenarios. Click TEST SETTINGS to verify that the database can be opened successfully.

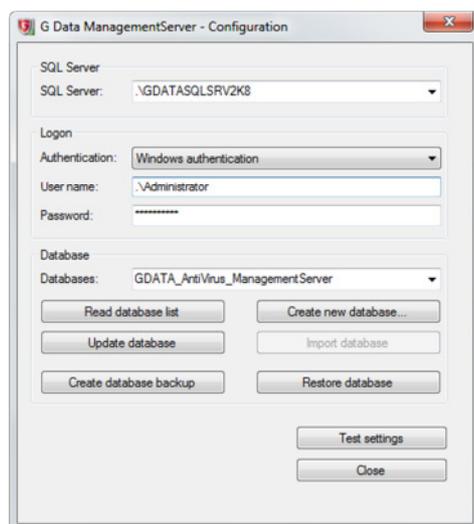


Image 17: GdmsConfig.exe

Click CREATE DATABASE BACKUP to select the folder. If a local database server was defined, a local folder can be selected. If a remote SQL server instance was defined, the path needs to be entered as a UNC path (for example, \\Backupserver\C\$\Backups). To restore a database backup, click RESTORE DATABASE. This will open a folder selection window which opens at the last used folder, but database backups can be restored from any folder.

GdmsConfig.exe can be used as a command line application, which helps in automating database backups. To make sure that the database is backed up regularly, a task can be added to the Windows Task Scheduler (Start, Run, *taskschd.msc*). A weekly task with the appropriate backup command makes sure that there is always a recent database backup to restore in case of emergency. The exact command depends on the SQL server deployment. If ManagementServer has been configured to use a local SQL Server Express instance, in most cases the command *gdmsconfig.exe /dbfullbackup /DBBackupFolder:folder* will do, replacing *folder* with an absolute path to the backup folder. It will connect to the database using the stored login information (which can be verified in the interface of GdmsConfig.exe), and store a backup in the specified target folder. This requires the local account to have sufficient permissions to connect to the database and to store the file in the target folder. If a database backup or restore folder had been previously selected in the GdmsConfig.exe interface, the */DBBackupFolder* parameter can be omitted to use the previously defined folder as the backup target. If ManagementServer is using a remote SQL Server instance, more parameters need to be defined. User name, password, database name and database instance should be included: *gdmsconfig.exe /dbfullbackup /DBBackupFolder:folder /Logon:username /Password:password /Database:database /ServiceInstance:sqlservername*. The backup folder needs to be defined as a UNC path and the logon account should have write permissions. For example, with database GDSQLSRV2K8 on server DATABASESRV, backup folder MMS on server BACKUPSRV, and user account BackupAdmin with password Password, the command looks as follows: *gdmsconfig.exe /dbfullbackup /DBBackupFolder:\\BACKUPSRV\MMS /Logon:BackupAdmin /Password:Password*

/Database:GDSQLSRV2K8 /ServiceInstance:DATABASERV. As with all backup-related tasks, it should be verified that the configured task runs successfully and indeed generates a backup in the desired location.

Using the same backup and restore functionalities, GdmmsConfig.exe can also be used in database tests and migration scenarios. See chapter 17.2 for more information.

4.7. Client deployment

Before commencing a network-wide client deployment, make sure that the default client security settings have been configured according to needs and policies. The CLIENT SETTINGS module of G Data Administrator can be used to configure security settings for all clients at once by selecting the main MMS, or by selecting one by one the groups for which separate settings are required. See chapter 8 for an overview of client security settings.

Clients need to be added to G Data ManagementServer's client management area before client software can be deployed to them. This allows administrators to keep an overview of network clients, even those to which software has not been deployed yet. To make MMS aware of clients, they need to be "enabled". This process can be initiated in several ways. G Data Administrator's SERVER SETUP WIZARD provides a list of network clients that has been detected in the local network, which can be used to enable one or more clients with a single mouse click. If a client is not listed, it can be enabled by manually entering its name or IP address. An alternative is the client management area. Right-click on a server or group and select ENABLE CLIENT. In the following window, you can enter any number of clients to be enabled. The third option is the FIND COMPUTER(S) window, which is located in the ORGANIZATION menu. A complete IP range can be searched for active network clients, which can then be enabled directly. Finally, Active Directory synchronization can automatically enable clients. Link an MMS group to an Active Directory Organizational Unit to automatically import its clients into MMS (see chapter 7.2). To enable clients manually by name, select the appropriate ManagementServer in the client management area, click ENABLE CLIENT on the toolbar and enter the client names.

Which clients to include in the G Data client deployment is up to the administrator. It is recommended to protect all machines in the company network. Leaving one of them out may provide an entry point for malware. All Windows, Linux and Android devices should be protected by G Data; this includes clients as well as servers. Although not all G Data security modules are as useful for servers (the client-oriented Firewall, for example), G Data Security Client's anti-malware modules, such as the file system monitor, provide excellent server protection as well. Note that, depending on the type of server, extra tests will need to be carried out. Stability and performance should be optimized. Several exceptions might have to be set for the file system monitor and scan jobs, such as regularly used database files on a database server, the e-mail database of a mail server, or several types of log and management files on a domain controller³.

Depending on the method of enabling clients, client installation may be triggered immediately. The Server setup wizard offers to initiate a remote installation on clients that it enables. When using Active Directory synchronization, remote installation can also be carried out directly. When the system requirements for remote installation have been met (see chapter 4.7.1), this is a very convenient way to make sure all clients are directly served with G Data Security Client. However,

³ Consult vendor documentation to find out about critical processes, files, and folders. Microsoft, for example, publishes guidance for its Exchange Server products at <http://technet.microsoft.com/bb332342>. Additional information is available in G Data TechPaper #0131 G Data Security Client & Microsoft Server.

there are several other installation methods, making sure that even clients to which the client software cannot or should not be remotely installed are secured quickly and easily. G Data Administrator can be used to build a client installation package. This executable file contains all of G Data Security Client's installation files and preconfigured settings to allow the Client to contact the MMS after installation. It can be copied to any network client and be run remotely or locally, completely unattended. Alternatively, a local installation is possible: the G Data installation medium also contains Security Client's installation files.

4.7.1. Remote installation

A remote installation of G Data Security Client can be initiated by the `SERVER SETUP WIZARD`, by Active Directory synchronization, or by selecting a client in the `CLIENTS` overview, right-clicking it and selecting `INSTALL G DATA SECURITY CLIENT`. A remote installation is the easiest way to install G Data Security Client and saves time, because the administrator does not need physical access to the client. However, a few configuration changes might be necessary in order to deploy G Data Security Client remotely:

- `Gdmms.exe` must be allowed network access (firewall exceptions need to be set).
- In a Windows Workgroup, `SIMPLE FILE SHARING` (Windows XP) or the `USE SHARING WIZARD` option (Windows Vista or Windows 7) must be disabled. It is enabled by default in all Windows client installations and can be disabled by opening any folder in Windows Explorer, clicking `Organize > Folder and search options > View`, and unchecking the option `USE SHARING WIZARD`.
- `USER ACCOUNT CONTROL (UAC)` must be disabled.
- Access permissions for the UNC shares `C$` and `Admin$` should be available. This can be a local or domain administrator account. An account password must be set: an empty password field is not permitted.
- The `REMOTE REGISTRY` service must be enabled in `SERVICES`.

On Windows 8, Windows 8.1, and Windows Server 2012 clients that are part of a workgroup, access to the UNC share `Admin$` is not available. To ensure a successful remote installation, a registry key needs to be added. Open Registry Editor and navigate to the key `HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\Policies\System`. Add a `DWORD` value named `LocalAccountTokenFilterPolicy` with value `1` to enable remote installations of Security Client. Alternatively, deploy Windows 8 workgroup clients using a local installation or a client installation package.

When all the requirements have been fulfilled, the remote installation(s) can be initiated. Depending on the way in which the installation was initiated, several questions may be asked. The language of the G Data Security Client installation can be chosen from a dropdown menu. Note that this setting cannot be changed afterwards: the client will have to be reinstalled if the language should be changed. The Firewall can also be installed remotely; if, for any reason, this component should not be installed at this time, it can be remotely installed later through the Firewall module (see chapter 13.1).

Behind the scenes, a remote installation can be carried out in two ways. If the clients meet the necessary prerequisites, the files are copied directly and entries are made in the registry. If the server can only access the hard drive and not the registry, or if other system prerequisites are not met, the entire setup program is copied to the client and started automatically at the next computer reboot.

If the latter option is not successful either, the installation is aborted.

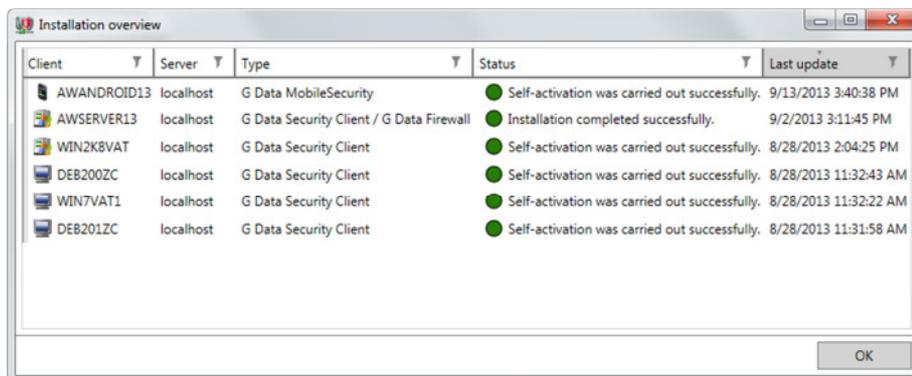


Image 18: G Data Administrator, Installation overview

To keep track of the progress of a remote installation, the INSTALLATION OVERVIEW window can be used. It opens automatically when a remote installation task is added, or can be opened by clicking the INSTALLATION OVERVIEW button in the client management area's toolbar. It lists all clients that have pending and completed remote installation tasks. The TYPE column shows the type of installation (G Data Security Client and/or G Data Firewall, or Subnet server). After a remote installation has been completed, the STATUS column will be updated. In some cases, such as when G Data Firewall is being installed, the client will need to be rebooted in order to complete the installation. The installation procedure will add a report to the REPORTS module if a reboot is required. If remote installation is not possible for any reason, the error will be displayed in the STATUS column. Further relevant error codes may be found in the registry (see chapter 4.7.4). If a remote installation fails repeatedly, G Data Security Client and G Data Firewall can be installed locally or using a client installation package.

4.7.2. Local installation

G Data Security Client can be installed locally on any client or server with a supported operating system. This is useful in scenarios where a network client is not in the same domain as the main MMS, the system requirements for remote installation cannot be met, or the client does not regularly connect to the network (laptops). The G Data installation medium contains an installer file, which can be run with local administrator rights on any client. The Security Client installer contains all available languages. The installation wizard is simple; all that is required is to enter the name of the main MMS that the client should connect to. After the installation has finished, the client will contact the MMS within minutes, after which it will be shown in the client management area.

Note that by default, the local installation wizard does not install the Firewall component. The recommended way to install the Firewall is remotely through G Data Administrator's Firewall module. If the Firewall should be installed locally as well, adding a parameter to the setup wizard enables its installation. Open a command prompt and navigate to the \Setup\Client folder on the G Data installation medium. Launch the installation wizard by typing `setup.exe /v"/INSTALLFW=yes"`. After completing the installation wizard, a reboot is required.

4.7.3. Client installation package

If a remote deployment is not possible, but local installations take too much time, the client installation package can be a practical solution. ManagementServer can create an executable which includes the latest version of the client program files and virus signatures, as well as preconfigured

settings to allow the client to automatically connect to the ManagementServer upon installation. The client installation package does not require any user interaction and is an ideal solution to quickly deploy G Data Security Client across a whole domain. Networks that make use of Active Directory can use startup scripts, so that the client retrieves the installation file upon login and executes it automatically in the background.

A client installation package can be created in G Data Administrator. Open the menu ORGANIZATION and choose the option CREATE G DATA SECURITY CLIENT INSTALL PACKAGE. Administrator will ask for the installation language and ManagementServer. Security Client will be installed in this language version and connect to the ManagementServer indicated here. Finally, a storage folder can be indicated. G Data Administrator will create the package in the background. The process can take a few minutes: the Administrator application should not be closed during that time. The client installation package includes the latest version of Security Client as well as the latest virus signature updates. This ensures that the client will be optimally protected immediately, without having to separately download updates from the ManagementServer. This does mean, however, that a new client installation package should be created regularly if the deployment is to last for a longer period of time.

As soon as the installation package has been compiled, it only has to be copied to the client(s) and run. Any existing Security Client installations have to be removed completely: the installation package will only work if no installation is present. Due to its file size, it is not recommended to run the client installation package from a network share: installation may fail. With about ten to fifteen minutes, the installation process takes slightly longer than a remote installation or a local installation from the installation medium. While Security Client is being installed, end users can continue making use of the client. After fifteen minutes, the client should be restarted twice to finish installation.

By default, the client installation package does not install the Firewall module. If required, the Firewall can be remotely installed afterwards through Administrator's FIREWALL module. Alternatively, the client installation package can be started from the command prompt with a specific parameter to force it to include the firewall in its installation procedure. On a command prompt, navigate to the folder that contains the package and type the following command: *gdclientpck.exe /fw*. This will launch the installation process that installs Security Client as well as its Firewall module. As with a Security Client-only installation, the Security Client plus Firewall installation will fail if an installation of Security Client is already present. It cannot be used to install only the Firewall module.

4.7.4. Troubleshooting

Whether the client is installed locally, remotely, or using a client installation package, complications can occur. Especially for remote installations it is difficult to directly spot installation errors, as the procedure is carried out remotely. If a client does not connect to the ManagementServer after being deployed, there can be several reasons. Most importantly, the client should be able to establish a network connection with the MMS, and vice versa (see chapter 4.8). If a network connection is available, but the client still is not connecting to the MMS, the installation procedure for G Data Security Client may have failed. If it fails, the installer saves an error code in the Registry. Using Windows Registry Editor, navigate to the key HKEY_LOCAL_MACHINE\SOFTWARE\Wow6432Node\G DATA\AVKClient⁴. The value InstallError will contain one of the following installation error codes:

⁴ On 32-bit systems: HKEY_LOCAL_MACHINE\SOFTWARE\G DATA\AVKClient.

Error code	Description
0 (Success)	Installation successful.
2 (Success)	The server name was successfully updated.
10	The provided password was wrong.
12	The path containing the source files does not exist.
13	No connection to Service Control Manager or Remote Registry.
14	Files could not be copied.
17	The client operating system must be NT-based.
18	Failed to connect to the Remote Registry service.
19	A client has already been installed.
22	Unable to read system and program path from the client.
23	Failed to connect to the Remote Registry service.
27	The machine does not have the required patch level (Service Pack).
28	Files are in use.
29	The client has not been rebooted after uninstalling a client.
40	Linux installation error: internal error.
42	Linux installation error: file not found.
44	Client not connected.
48	Authentication failed (identical to error code 10).
49	Firewall installation error.
50	An attempt was made to install the firewall without a client.
51	The firewall has already been installed.
52	An attempt was made to install the firewall on a server operating system.

Most of the error messages can be resolved by ensuring that the correct permissions are set and that the system meets the requirements for remote installation (chapter 4.7.1). If an installation error occurs, it should be made sure that no Security Client files are left on the client before reattempting to install it. The AV-Cleaner tool can help clean up installation leftovers (chapter 17.2.4).

4.7.5. Mobile clients

Android client installations can be initiated from G Data Administrator. There are a few prerequisite settings. On the tab `SERVER SETTINGS > MOBILE`, enter a password under `AUTHENTICATION FOR MOBILE CLIENTS`. The password will have to be entered in the app to allow it to authenticate with G Data ManagementServer. In addition, make sure that the ManagementServer can accept connections on TCP port 7184. Mobile clients connect to the server on this port to download the app's installation file.

The deployment process is carried out via e-mail. In the client management area, click `SEND INSTALLATION LINK TO MOBILE CLIENTS`. A list of e-mail addresses can be entered: an activation e-mail can be sent to any e-mail address. When opened on the client, a link to download the Internet Security for Android installation file from the ManagementServer will be included. Tap the download link to download the installer APK file. Note that the option `UNKNOWN SOURCES (ALLOW INSTALLATION OF NON-MARKET APPS)` needs to be enabled in order to install APK files. This option is usually found in Android's system

menu **SETTINGS > SECURITY > DEVICE ADMINISTRATION**. After opening the APK file and confirming its requested permissions, G Data Internet Security will be installed and can be started from the Android app menu.

Through the Settings icon in the top right corner of the screen, the app can be configured to allow remote administration. Tick the checkbox **ALLOW REMOTE ADMINISTRATION** and enter the name or IP address of the ManagementServer under **Server address**. Under **DEVICE NAME** you can enter a name that will be used to identify the device in G Data Administrator. **PASSWORD** should contain the password that you entered in G Data Administrator (which is also listed in the installation e-mail). The device will be listed among the other clients in G Data Administrator's **CLIENTS** module and can be managed from there. If it does not appear automatically, reboot the device to force it to check in with the ManagementServer.

4.7.6. Linux clients

Like their Windows counterparts, Linux clients can be linked with the G Data ManagementServer infrastructure, centrally managed via G Data Administrator and supplied with signature updates. There are two types of Linux clients. The Workstation client provides a file system monitor with a graphical user interface and functionality similar to the AntiVirus module for Windows. The Server module is meant for Linux servers that communicate with Windows network clients (using the SMB protocol). This module carries out a file scan on every access, so malware cannot migrate from the server to the Windows clients (or vice versa).

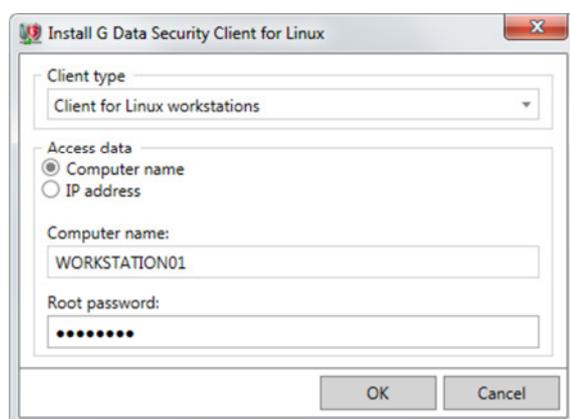


Image 19: G Data Administrator, Install G Data Security Client for Linux

To carry out a remote installation using G Data Administrator, select any Linux client in the **CLIENTS** module. From the menu **CLIENTS**, choose the option **INSTALL G DATA SECURITY CLIENT FOR LINUX**. Select the type of client and enter the computer name (if it has been registered with the network's DNS server), or its IP address. To install Security Client for Linux, the root password has to be entered. A root password must be set, even if the client's Linux distribution does not do so by default (such as Ubuntu). Click **OK** to initiate the remote installation. The **INSTALLATION OVERVIEW** window will show the installation's progress.

If for any reason a remote installation is not possible, Security Client for Linux can be installed locally as well. Insert the G Data installation medium and navigate to the folder **\Setup\LinuxClient**. The folder contains three files: an installer, an uninstaller, and virus signatures. Run **installer.bin** to install the software. The signatures do not need to be installed, as Security Client will automatically obtain the latest ones from the server after the installation.

4.8. Finalizing deployment

After completing the deployment of both server(s) and clients, it is important to verify that all processes are running correctly and that all security measures are in place. Most importantly, all clients should be able to connect to the MMS. By default, each client reports its status to the MMS every five minutes (except during a scheduled scan). The CLIENTS module of G Data Administrator will help locate and troubleshoot connection problems. Firstly, check if all network clients are listed in the client selection area. If a client is missing, it may not have been deployed correctly. Depending on the method of client deployment, try to add the client manually and enable it, or check the network's Domain Controller to see if the client has been added to Active Directory. If the client is listed, use the LAST ACCESS column to find out when the client last connected to the MMS. Make sure that the client is connected to the network and turned on. TCP traffic must be allowed through the relevant ports, both on the client (7167/7169) and the server (7161). Finally, the client must be able to resolve the IP address of the server. To test the connection, use the *telnet* command on the client to connect to the server: *telnet <MMS IP> <MMS port>*. If the client can connect to the server, an array of cryptic characters will be displayed on the prompt. If there is no connection, an empty input window appears. See chapter 7 for more information about the CLIENTS module and client management.



Image 20: Telnet (connected to MMS)

With all clients regularly connecting to the MMS and protected by the default network settings, the core deployment has been finished. However, there are more services to configure. To allow for quick configuration changes, configure the G Data solution for remote access (see chapter 5). Customize real time protection settings for each client as appropriate (see chapter 8). Recurring tasks can be scheduled to perform malware scans (see chapter 9) and backups (see chapter 11). If the deployment scenario includes the Firewall component of the G Data solution, make sure to configure and customize its settings (see chapter 13). Most settings can be configured using G Data Administrator, but some advanced settings can be accessed by editing configuration files or using specialized G Data configuration tools (see chapter 17).

5. Remote administration

Some configuration changes need to be carried out unexpectedly. It can be necessary to gain access to G Data ManagementServer's configuration from a machine without configuration tool G Data Administrator, or while on the road. G Data offers full configuration possibilities through the browser, as well as a selection of the most-used options for mobile devices (such as smartphones and tablets).

During the installation of G Data ManagementServer, configuration tool G Data Administrator will be installed on the same machine. G Data ManagementServer can then be configured by physically using the server machine or by logging in to it using Windows' Remote Desktop Protocol, or any third-party remote control solution. Additionally, G Data Administrator can be run from other machines without the need to open a session on the server, by installing it on any machine with network access to the server.

As an optional part of the deployment process, the configuration capabilities of the G Data solution can be made remotely accessible. G Data Administrator can be installed and configured to be accessed from outside the network, but if installing the Administrator software is not an option, G Data WebAdministrator offers a browser-based interface that provides access to all settings and modules. For mobile users, G Data MobileAdministrator is the perfect interface, offering the most commonly executed tasks, such as client and security management, and checking reports. For ManagementServer to be administered remotely, TCP port 7182 needs to be accessible (see chapter 2.2).

5.1. Desktop application

The default deployment of G Data ManagementServer installs G Data Administrator on the same machine. With physical or remote desktop access to the server, administrators can log in to G Data Administrator to get access to all modules. In cases where desktop application access to the server is not possible or not practical, G Data Administrator can be installed on any other Windows client, as long as it can reach the MMS.



Image 21: G Data Administrator, Login

Use the G Data installation medium to install G Data Administrator on the PC from which configuration tasks will be carried out. When logging in, enter the IP address or (if resolvable) the name of the MMS server machine as SERVER address. Make sure that the server port (default: TCP 7182) is not being blocked by the firewall, and forward it on router level if necessary.

5.2. Browser

Taking the time to install G Data Administrator on a machine is not always an option. Local policies can prevent software from being installed, or an urgent issue requires immediate attention and leaves no time for a software installation. In these cases, it is very practical to be able to configure G Data ManagementServer using only a browser. The web-based module G Data WebAdministrator offers this possibility. Most commonly, WebAdministrator is deployed to an existing web server in the enterprise network, but it can be installed to any Windows machine that is running Microsoft Internet Information Services (IIS). The following versions of IIS are supported, with their respective operating systems:

IIS version	Operating system
5.1	Windows XP Professional
6.0	Windows Server 2003
7.0	Windows Server 2008, Windows Vista
7.5	Windows Server 2008 R2, Windows 7
8.0	Windows Server 2012, Windows 8

Microsoft IIS needs to be installed before WebAdministrator can be deployed. Each Windows version listed above includes the IIS component, but often it needs to be enabled manually. For Windows Vista, Windows 7 and Windows 8, open the WINDOWS FEATURES panel (found under Control Panel > Programs and Features). Select INTERNET INFORMATION SERVICES to install the complete web server package, or pick individual components. Also enable IIS 6 MANAGEMENT COMPATIBILITY > IIS METABASE AND IIS 6 CONFIGURATION COMPATIBILITY, as WebAdministrator depends on it. Click OK to install IIS and restart the machine if prompted to do so.

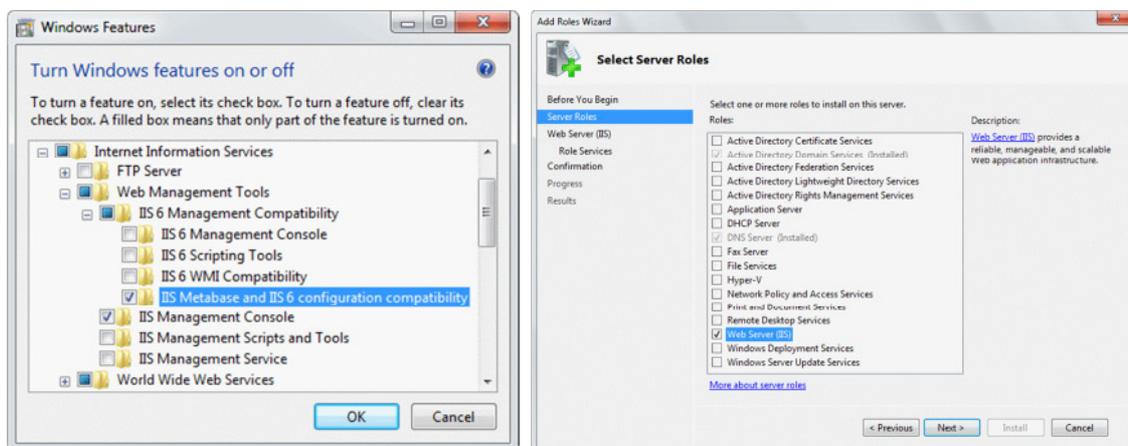


Image 22; 23: Windows 7, Windows Features; Windows Server 2008 R2, Add Roles Wizard

Using Windows Server 2003, start the Manage Your Server application from the Start menu. In Windows Server 2008 and 2012, this function has been renamed Server Manager. Both applications feature the possibility to add Roles to the current server configuration. In Windows Server 2003, the

appropriate role is called APPLICATION SERVER (IIS, ASP.NET); in Windows Server 2008 and 2012, WEB SERVER (IIS). For the latter two, IIS 6 METABASE COMPATIBILITY needs to be selected on the ROLE SERVICES panel. After installing the web server role (and possibly restarting the server), verify that the web server is accessible by opening `http://localhost` in the local browser.

As with any website, accessing G Data WebAdministrator through the browser can expose HTTP traffic to attackers with network access. Especially in scenarios where G Data WebAdministrator will be accessed from outside the enterprise network, securing the traffic is recommended. This can be done using an SSL certificate. Certificates are available for purchase from Certificate Authorities (CAs) or can be generated locally for free and be self-signed. The former option is recommended for cases where WebAdministrator will be accessed from outside the enterprise network, but will incur additional costs if the enterprise does not already own one or more certificates. The latter option can be configured easily, and will protect against casual eavesdropping on the HTTP traffic, but is more vulnerable to a man-in-the-middle attack.

Using Windows XP Professional or Windows Server 2003, an SSL certificate can be added by using the free Microsoft tool SelfSSL, available from the Microsoft website as part of the IIS 6.0 Resource Kit Tools⁵. After installation, open the SelfSSL command prompt through Start > Programs > IIS Resources > SelfSSL. A self-signed certificate can be assigned to the local website by entering a single command: `selfssl /N:CN=localhost /K:2048 /V:365 /S:1 /T`. Confirm the certificate creation by pressing Y. This will create a certificate for the default IIS site on the local server, and add `localhost` to the list of trusted certificates. The key length will be 2048 and the certificate will be valid for 365 days. If the site is not the default site of IIS, look up its IDENTIFIER in Start > Administrative Tools > Internet Information Services (IIS) Manager and change the parameter `/S:1` accordingly.

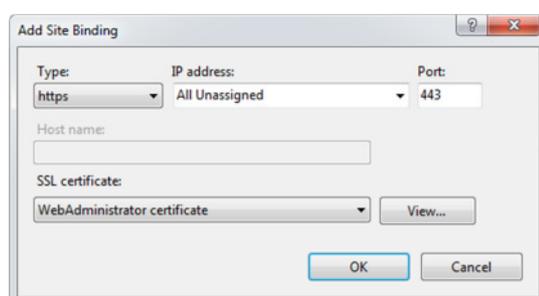


Image 24: Internet Information Service (IIS) Manager, Add Site Binding

Using Windows Vista, Windows 7, Windows 8, Windows Server 2008/R2, or Windows Server 2012, open Internet Information Services (IIS) Manager by clicking Start > Run (or, alternatively, by holding Windows-key + R) and entering the command `inetmgr`. Select the local web server in the Connections panel. In the middle of the screen, navigate to the IIS category and double click on SERVER CERTIFICATES. On the ACTIONS panel, click CREATE SELF-SIGNED CERTIFICATE. After entering a friendly name for the certificate, it will be created and listed in the SERVER CERTIFICATES panel. Note that the default expiration date of the certificate is exactly one year ahead of the date of creation. To apply the certificate to site communication, select the appropriate site in the CONNECTIONS panel. On the ACTIONS panel on the right, choose BINDINGS. Click Add to add a new binding. Select `https` as type and select the new certificate in the SSL CERTIFICATE dropdown. Click OK to add the binding.

With IIS configured, G Data WebAdministrator can now be installed. Use the setup wizard on the G Data installation medium to install WebAdministrator. Microsoft .NET Framework will automatically be installed if the server does not yet have the required version. After installation,

⁵ See www.microsoft.com/en-us/download/details.aspx?id=17275.

WebAdministrator will be accessible in the browser by opening the subfolder /GDAdmin, such as *https://10.0.2.150/GDAdmin* (or *http://* if no SSL certificate has been installed on the web server). The folder will be different if the installation folder has been altered. Because of the self-signed certificate, browsers may issue a warning before opening WebAdministrator. The communication, however, will still be fully encrypted. If the Silverlight browser plugin has not yet been installed, the user will be prompted to do so upon the first visit.

G Data WebAdministrator can be used to log in to any ManagementServer. Its login methods, interface and functions are identical to those of G Data Administrator. Any configuration and management tasks can be carried out through the web interface.

5.3. Mobile

For configuration tasks that need to be carried out right away, G Data Administrator and G Data WebAdministrator are not always the perfect solution. For cases where no software or desktop browser access is possible, G Data has developed MobileAdministrator. It offers access to the most commonly used functions of G Data Administrator in a mobile-optimized web interface.

MobileAdministrator can be used on all smartphone platforms and on all tablets and does not require the Silverlight plugin, unlike WebAdministrator. The web application can be used to manage clients and to stay up to date with the latest reports about malware infections, PolicyManager requests and more. It offers effective client management and reports. The web application does not just provide passive reporting capabilities, but supports direct responses. Malware infections can be checked and directly acted upon. Files can be quarantined or moved back and PolicyManager reports used to directly edit white- or blacklists. The web application can also be used to quickly gain an overview of the status of all network clients. Reports can be defined and previewed using the mobile ReportManager module.

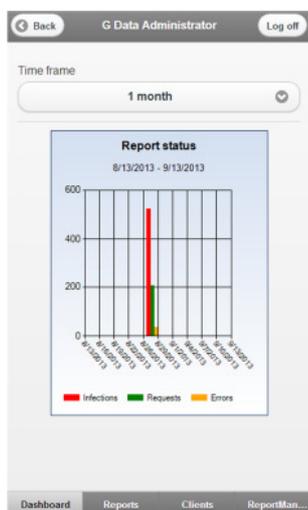


Image 25: G Data MobileAdministrator, Dashboard

MobileAdministrator is, like WebAdministrator, a web application. It can be installed from the G Data installation medium, on top of Microsoft Internet Information Services (IIS). MobileAdministrator requires at least Windows 7 or Windows Server 2008 R2. For more information about configuring IIS, including an SSL certificate, see chapter 5.2.

5.4. MasterAdmin

Although G Data Administrator, WebAdministrator and MobileAdministrator can be used to log in to any ManagementServer, effective management of very large networks should be carried out with MasterAdmin. This version of G Data Administrator allows management of multiple MMS within the same interface, streamlining configuration and deployment. To manage multiple servers, MasterAdmin functionality can be enabled in G Data Administrator. ManagedService partners, as well as end customers who are managing a large network with multiple MMS installations, can request a MasterAdmin activation code from G Data. On the regular login screen, select **MANAGE MULTIPLE SERVERS** to enable the appropriate login options. Enter the activation code and a username and password of choice. After successfully logging in, the **MASTERADMIN WIZARD** will be started automatically. Using the wizard, the management servers that will be administered remotely can be added. Enter the server's domain name or IP address and its user name and password. To tell multiple servers apart in the MasterAdmin interface, enter an alias name. Click **NEXT** to add a new server or **FINISH** to close the wizard. The MasterAdmin wizard can be opened at any time from the **ADMIN** menu.

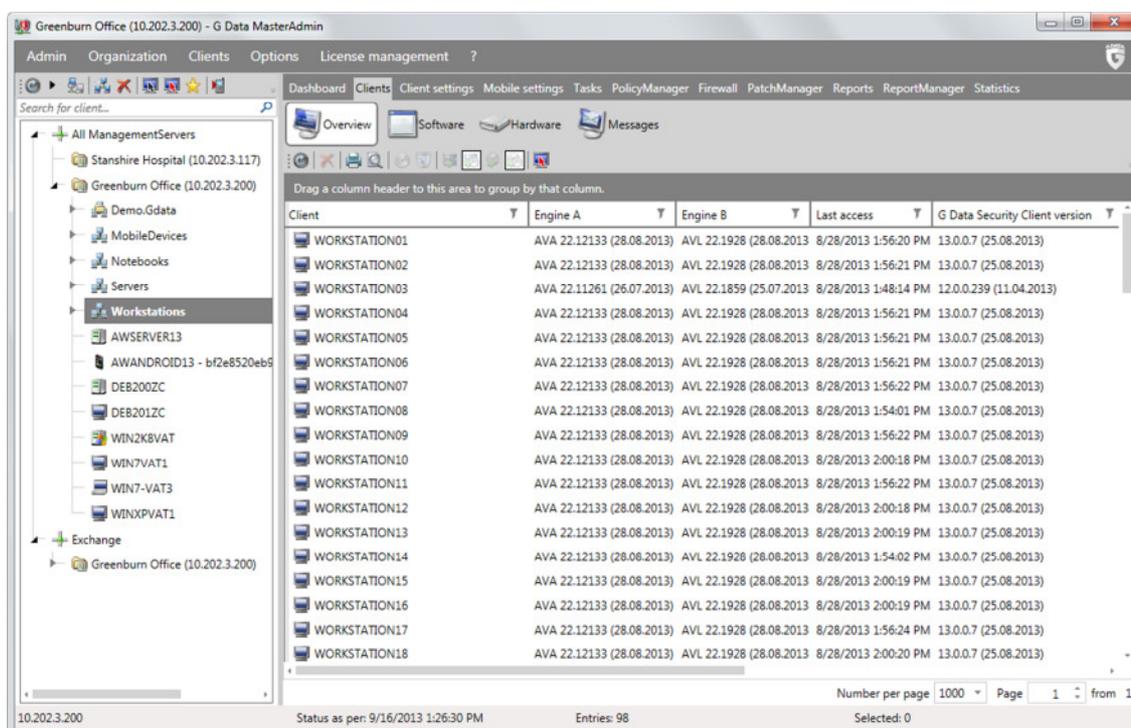


Image 26: G Data MasterAdmin

After adding servers, MasterAdmin's options are virtually indistinguishable from G Data Administrator's regular functionality. Each MMS and its clients can be managed by selecting it in the client management area. Depending on the selected server's license, the appropriate modules will be shown on the right.

Section B: Using G Data business solutions

6. Dashboard and monitoring

After successfully finishing the deployment of G Data business solutions across the network, all clients will be continuously protected. However, each client's security status needs to be checked regularly. For various reasons, protection may be interrupted. Network problems or a lack of disk space can prevent updates from being distributed. Program or driver updates may interfere with a protection module or performance may be impacted. At any rate, running a security solution deployment is not a self-sufficient process. Ensure that all relevant information can be accessed efficiently and that administrators are notified of possible service interruptions.

G Data security solutions offer several notification possibilities. At-a-glance statistics can be found in the `DASHBOARD` module. It provides charts with information about infections, client connections and more. Common tasks can be carried out directly, such as updating virus signatures or enabling the firewall. A more comprehensive view of the statistics can be found in the dedicated `STATISTICS` module. Individual clients send reports about malware infections, corrupted files, and more. Reports can be accessed through the module of the same name, where direct action can be taken for some messages. Finally, the `REPORTMANAGER` module allows administrators to compile their own reports from several modules and have the results e-mailed regularly.

As a general recommendation, administrators should try to set up a balanced notification system. By combining statistical information, e-mail notifications and ReportManager reports, a good overview can be obtained. Beware of information overload: by having the software report more than necessary, important notifications can slip by unnoticed.

6.1. Statistics

By default, opening G Data Administrator leads directly to the `DASHBOARD` view. It allows the administrator to get an overview at a glance of the key statistics. While the `DASHBOARD` view itself does not generate e-mail reports (those can be managed through the `ALARMS` setting and the `REPORTMANAGER` module), the charts are a very valuable tool to spot anomalies in the network client protection status.

The `G DATA SECURITY STATUS` column offers an indispensable overview in numbers. It is the first panel that should be checked upon launching G Data Administrator. If one of the managed systems does not have G Data Security Client running, the Security Status column will show the affected client. Similarly, clients with outdated virus signatures will be listed, allowing administrators to update them immediately. If the Monitor, E-mail checking, OutbreakShield, Web protection, BankGuard, or Firewall components are disabled, they can be enabled directly. If the Firewall has not been installed on certain clients, a remote installation can be initiated.

For machines where G Data Security Client has been installed, the `CLIENT CONNECTIONS` overview shows the last connection. By keeping an eye on the `CLIENT CONNECTIONS` chart, administrators can check when machines have last connected to the MMS and spot potential server or network problems. Machines that do not connect to the MMS regularly do not receive virus signature updates and will not be able to synchronize tasks and other settings. Some types of machines, such as laptops for external personnel, may be without access while on the road for longer periods of times. Regular desktop clients, on the other hand, can typically connect to the MMS more often, and if a considerable

percentage of clients has not been able to connect for more than 3 or 7 days, server or network problems may be occurring. If a machine is connected to the network and turned on, the connection between client and MMS should be investigated (see chapter 4.8), as well as network-level protection and port configuration (see chapter 2.2).

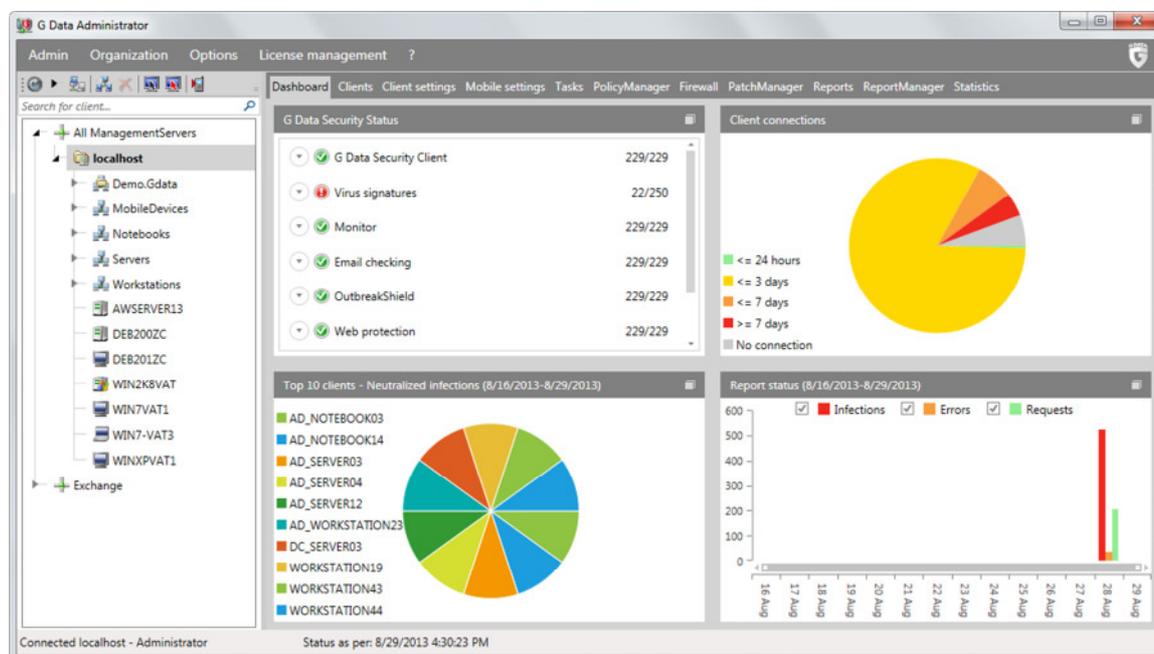


Image 27: G Data Administrator, Dashboard

The TOP 10 CLIENTS – NEUTRALIZED INFECTIONS overview helps located problematic clients. The pie chart shows the clients which produced the most virus reports. While G Data successfully repels these infections, the fact that particular clients are often hit by malware attacks can indicate problems. One of the other protection mechanisms may be configured incorrectly, or the end user may be particularly liable to be attacked, either because of a targeted malware attack, or because of careless (browsing) habits. Checking the security configuration for that client is recommended. If end user (mis)behavior is a likely cause, setting PolicyManager policies can help restrict access to dubious resources (see chapter 14). By default, the pie chart shows repelled infections from the last two weeks only. By clicking on the calendar icon in the top right corner, the period can be configured to cover the last few days, weeks or months, or a manually defined period.

REPORT STATUS shows an overview of reports (accessible through the Reports module). Infections, errors and Firewall and PolicyManager requests are laid out in a chart. Excessive errors from one of the modules or other notable spikes in the chart can be researched by opening the Reports module and investigating the individual reports. Like the NEUTRALIZED INFECTIONS chart, REPORT STATUS can be configured to show information for only a certain period of time.

In addition to the DASHBOARD, more extensive information can be gathered from the STATISTICS tab. On different panels it provides an overview of the protection status of the network. The CLIENTS panel displays the number of clients that have monitor, OutbreakShield and e-mail protection enabled, as well as the current status of engine and settings, and more. The CLIENTS panel can be configured to display up to eight types of statistical information, either as table (text), bar chart or pie chart. DETECTION METHOD displays the components that detected malware, information that can help in determining malware attack vectors, as well as identifying underperforming or misconfigured protection components. The VIRUS HIT LIST panel identifies the malware that has been detected most

often, information that can also assist in analyzing attacks. Finally, the HIT LIST OF NEUTRALIZED INFECTIONS shows the clients that have most often been targeted by malware. When one or more clients stand out, further inspection and active protection might be necessary, either because of a (semi-)targeted attack, or because of carelessness on the side of the end user. Using the PRINT function, statistics can be printed, for example to be used in external reports.

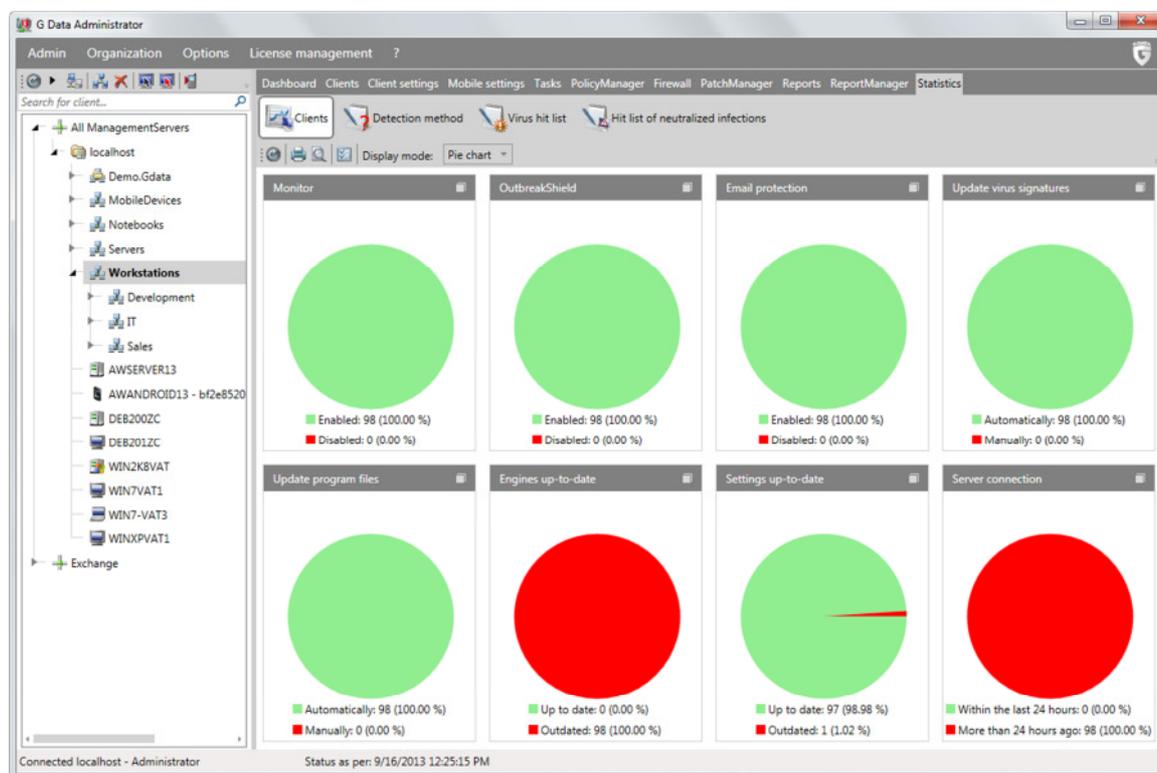


Image 28: G Data Administrator, Statistics

6.2. Reports

The REPORTS module is the place where reports (notifications) from all security modules can be found. While threats are blocked fully automatically, it is essential to keep up with reports. They provide detailed information about the status of network clients, potential security problems and other modules. For example, when using the POLICYMANAGER or PATCHMANAGER modules, reports need to be monitored for policy change requests or patch deployment or rollback requests. Because of the sheer amount of reports that are potentially produced, it is essential that administrators know how to filter reports in order to reduce the signal to noise ratio.

Reports are generated by the different modules of the G Data security solution. By default, all reports are listed in a flat list, in reverse chronological order (newest first). Depending on the number of clients, this list can seem overwhelming. Use the list control at the bottom of the window to set the number of items displayed per page. The group bar above the column headers can be used to group reports by column. For example, dragging the REPORTED BY column to the group bar will display all reports, grouped by the module they were reported by. Furthermore, the REPORTS module offers a vast range of options to filter the list. Using the toolbar, different types of reports can be hidden, such as reports that depend on other reports on the list (preventing duplicates), or reports that have already been read. It is also possible to view reports of specific categories, such as viruses that have not yet been removed, quarantine contents, or BankGuard reports. Using the TIME FRAME option

allows administrators to limit the number of displayed reports to a specific period of time.

The most common security modules producing reports are Monitor, Scanner and BankGuard. Most reports that are being produced concern malware that has already been blocked. Nevertheless, this type of reports is important. It allows administrators to gain insight in the state of their network and which machines have been compromised. Furthermore, false positives can be identified and added to an exception list. Monitor and Scanner reports show the action that was taken when the virus was found, such as virus removal, file quarantine or file deletion. Files that have been quarantined can be moved back (in case of a false positive) or cleaned directly from the Reports module by clicking the appropriate toolbar button. This can be risky: if a file has not been completely cleaned, it will remain infectious and can do further damage to the client system. The Reports module can also be used to directly define an exception. If a file is unjustly identified as malware, right-click on the report and select Define monitor exception to mark the file as safe for Monitor scans. Clients that have the Firewall module installed will also generate reports, for example when an application should be unblocked. Through the Properties option of these reports, firewall rules can directly be amended to allow the application.

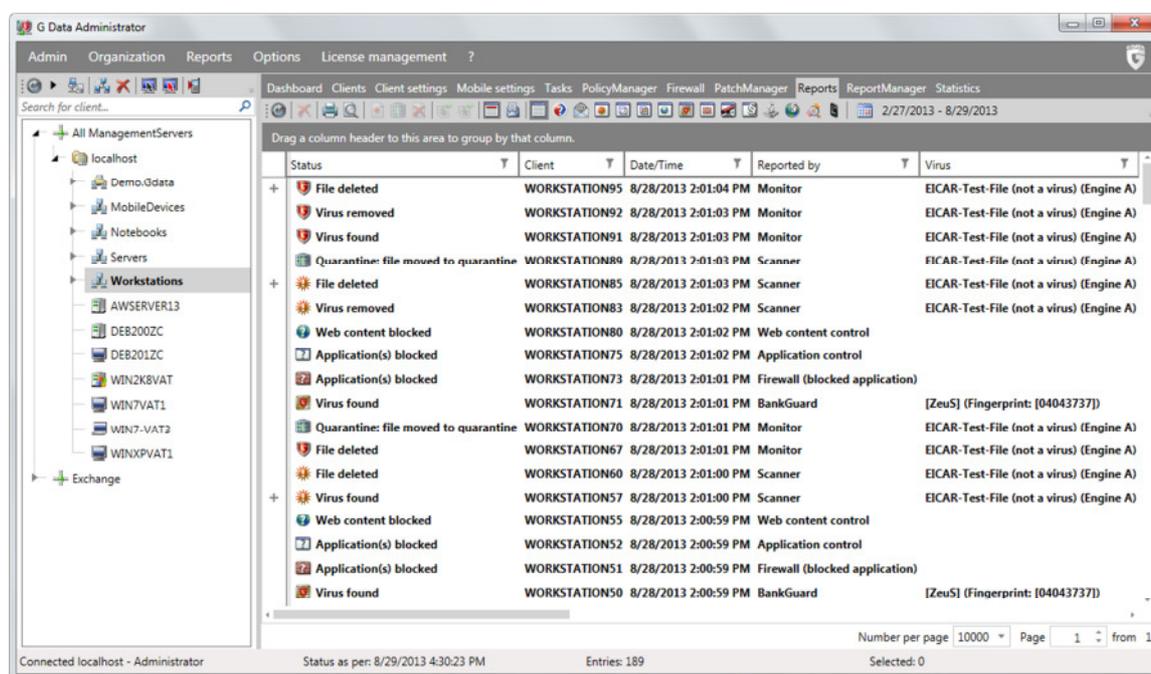


Image 29: G Data Administrator, Reports

In addition to security reports, some supporting modules also generate reports. The PolicyManager module logs web content, applications and devices that have been blocked due to policy settings and for which the user has requested access. By opening the report's properties, access rights can be added to the respective whitelist or blacklist. Similarly, PatchManager produces reports that can contain user requests, such as a patch distribution or rollback request. The Internet Security for Android component adds a report to the module when an end user requests access to a phone number that has been previously blacklisted.

Reports can be cleaned up automatically. The SETTINGS tab of the SERVER SETTINGS window allows administrators to specify automatic report deletion for logs that are older than a certain number of months. This setting can drastically clean up reports, but may remove vital notifications if they have gone unread for a longer time. It is recommended that the report cleanup setting is only enabled as automated cleanup measure in a workflow where reports are timely read (or distributed via e-mail).

It is strongly recommended to configure **ALARMS**. This setting allows administrators to have critical reports sent directly to their e-mail address, such as reports about virus detection, outdated clients, firewall actions and PolicyManager requests. While not every virus event is critical, alarms are key in staying updated with what happens on the network. They can be used to notify administrators or their on-call teams in case of emergency, or to provide information to users that do not have access to G Data Administrator or do not often log in. Use the **OPTIONS > ALARMS** window to configure which reports should be sent to whom. Add a recipient group containing at least the e-mail addresses of the administrator(s) and IT personnel providing emergency response⁶. By default, only outdated virus signatures and problems with client virus signature databases or program files are reported. These notifications typically require (semi-)immediate responses, because they directly affect client security. Reports on virus detection and applications blocked by the firewall are of informational nature, but can be enabled to gain a more complete overview. When using the PolicyManager module, enable **POLICYMANAGER REQUESTS** notifications to receive an e-mail when users file requests through PolicyManager (such as the whitelisting of certain websites or applications).

6.3. ReportManager

ReportManager allows administrators to compile reports from different customizable information modules. It offers a highly customizable experience, and a deep insight into network and security status. ReportManager can be seen as the proactive counterpart to e-mail notifications. With e-mail notifications set for events such as virus alerts and PolicyManager requests, administrators can quickly react to events. ReportManager, on the other hand, offers analytical information and lets administrators plan ahead. As a complement to regularly logging in to G Data Administrator, regular e-mail reports should be a part of the day-to-day administration workflow.

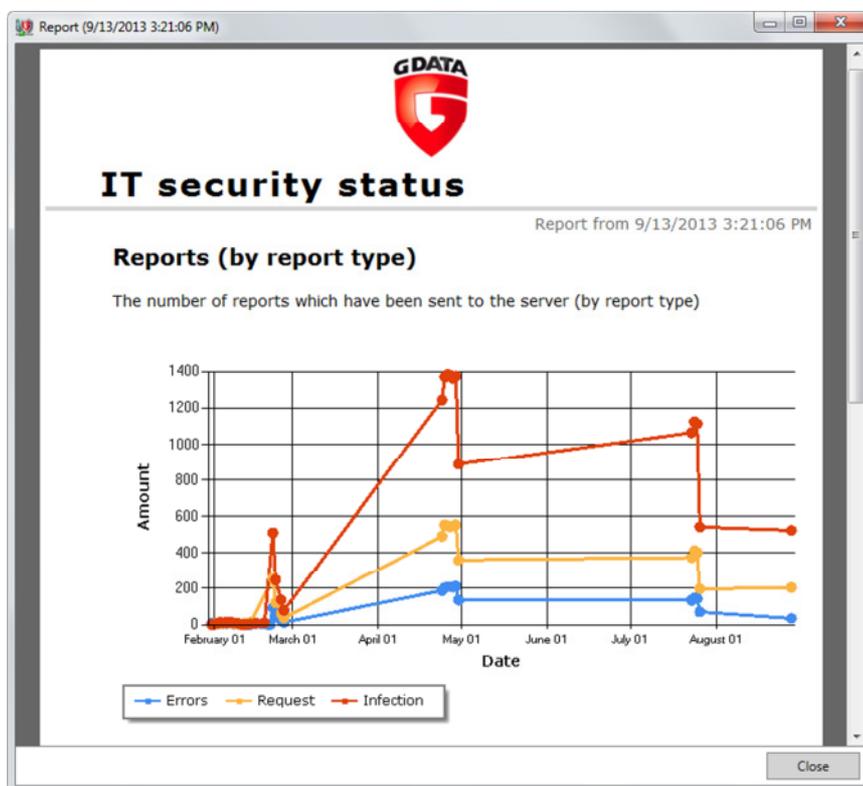


Image 30: G Data Administrator, ReportManager, Report

⁶ Make sure that e-mail recipient groups and an SMTP server have been entered in the Email settings window. See chapter 4.4.

The ReportManager module allows reports to be created containing information from different parts of the G Data security solution. Charts can be combined into reports to create information tailored to specific audiences. For example, a report for management personnel would contain high-level statistics about the number of protected clients, up-to-date software versions and patch status. Technical and administrative personnel could receive detailed statistical information about network clients and repelled virus infections.

Like scan or backup jobs, ReportManager can schedule report generation. Upon defining a new report, it can be configured to be run once, daily, weekly, monthly, quarterly, biannually or annually. While a one-time report can be a great way to gain insight into a specific component of G Data protection, the power of the ReportManager module is the possibility to automatically provide status reports over a longer period of time. Reports can be used as source of actionable information, helping administrators to stay up to date with what is happening within the network, but can also be saved for later external reporting.

It is recommended to configure reports to be e-mailed to the relevant personnel regularly. Use e-mail recipient groups to coordinate report delivery (see Alarm configuration in chapter 6.2). The selection of modules depends on the target audience and information needs. The CLIENT GENERAL and CLIENT PROTECTION categories feature a number of charts with concise, concrete information. As a source of actionable information, it is recommended to create at least two types of daily reports: a management version of the report with high-level information, as well as a technically detailed status report for administrators. When using reports as a means of chronicling network protection over a longer time, a (bi-)weekly or monthly report can be defined. It is not necessary to define e-mail recipients, for example if the preferred method of reading reports is opening them through G Data Administrator. When using the PatchManager module, the additional category PATCHMANAGER features several charts that can be added to reports chronicling the status of patch deployment. This is especially helpful when setting up a patch management procedure (see chapter 15).

While it can be tempting to rely on daily or weekly reports alone, they are not a replacement for regular checkups, such as logging in to G Data Administrator and checking the DASHBOARD and REPORTS sections.

7. Managing clients

Proper network security starts with proper client management. By keeping tracks of active network clients and organizing them into relevant groups, administrators can reduce the chance of complications when deploying software and configuration changes. Ideally, the layout of every enterprise network has been defined in advance, following logical structures (see chapter 1). These structures can be mirrored in the G Data security solution to allow for more efficient client management. Through the use of client groups, administrators can create logical entities that can be managed all at once. Security settings, backup jobs, policies and many more options can be applied to groups. When using Active Directory, entities can automatically be imported to save time in configuring groups. For networks that are big enough to have deployed Active Directory previously, it is recommended to integrate AD entities into G Data Administrator groups. This reduces the time needed to reorganize groups, and cuts down on the amount of time needed for the security configuration of newly deployed clients.

7.1. Using groups

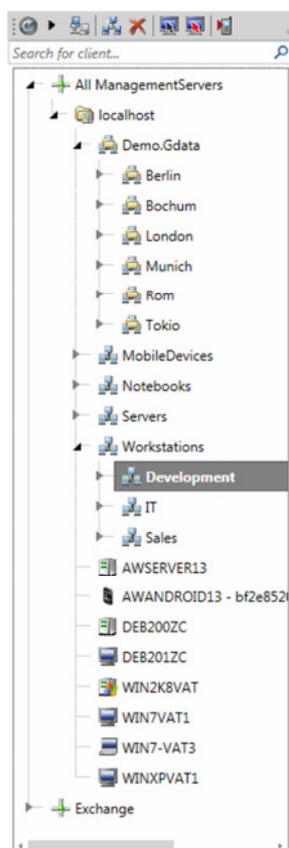


Image 31: G Data Administrator, client management area

The client management area of G Data Administrator offers some tools to divide clients into groups. Groups can be made and populated manually, or automatically using an existing Active Directory setup (see chapter 7.2). In its essence, groups function like a folder structure. Clients can be dragged freely between groups, allowing administrators to quickly create elaborate group structures. However, some planning should go into creating groups. There are several reasons for grouping clients together. Machines that are being used for similar tasks (e.g. development, back office, sales) will have similar software deployments. By grouping them together, identical software use policies can be rolled out. When using PatchManager, each similarly deployed group can be tested quickly by deploying patches to one designated testing client. A second possibility, partially overlapping with task- or software deployment-based grouping, is physical grouping. Client machines that are physically located close to each other can be bundled. Often this corresponds to branch offices or departments (in which case there is an overlap with task-based grouping). Location-based grouping has the advantage that it is easier to plan maintenance and configuration changes that require physical access to the machines. Many of these considerations are identical to the questions that need to be asked during the definition and setup of the local network. A simple solution is simply mirroring the network layout (see chapter 1.1).

Having defined groups, administrators can take advantage of several options. The most basic way to use groups is to apply settings to multiple clients at once. Select a group in the client management area and edit its settings (for example, protection settings in the CLIENT SETTINGS module).

Clicking APPLY will then apply the chosen settings to all clients in the group.

Alternatively, settings from one client can easily be transposed to the whole group by clicking APPLY TO ENTIRE GROUP. This makes it very easy to experiment with different settings on one client and afterwards applying the tried-and-tested configuration to the rest of the group.

Groups make client selection significantly easier. Processes like patch deployment or the planning of

staged update distribution can be carried out swiftly by selecting the appropriate group. Giving each group a descriptive name is essential, corresponding with the physical location, task type, department or any other logical entity that the system was based on. Additionally, settings can be easily exported and imported. By selecting a client, right-clicking and choosing `EXPORT SETTINGS`, the settings from the PolicyManager and Client settings modules can be exported to a `.dbdat` file. `IMPORT SETTINGS` can be used to selectively import settings back from a `.dbdat` file to a client or group.

7.2. Integrating Active Directory

Active Directory is a directory service for Windows networks. One or more AD servers within the network host information about units within the network, such as users, computers, and resources. Active Directory is a great help in organizing a network, authenticating users and coordinating access to resources. Effort that has been put into setting up a network structure in AD does not have to be duplicated when deploying and using a G Data security solution.

Any AD container or organizational unit (OU) can be linked to a group in the client management area. All current and future computers that are contained in the item will be automatically added to the overview in G Data Administrator. Create a group and right-click it to assign an AD item to the group. Each group can only be assigned one AD item. The dialog window allows administrators to select items from the default domain, or to add another domain by entering domain controller, domain, username and password. Upon confirming the AD group and refreshing the client management area view, the group is transformed into an AD-linked group and will have all computers from the AD unit added to it. G Data ManagementServer automatically synchronizes AD items with the domain controller every 6 hours, a value which can be changed in the Server settings window. AD-linked groups can also be updated directly by right-clicking them (or their parent MMS) and choosing `UPDATE ACTIVE DIRECTORY`.

Assigning AD items to groups reduces the time needed to manually add clients to groups. For new clients, configuration settings and other policies can be inherited automatically from the parent group, eliminating the need to configure new machines immediately. Further automation can be realized by automatically installing G Data Security Client on machines that are newly added to the AD group. This option can be enabled while assigning an AD item to a group in G Data Administrator. On the condition that the system requirements for a remote installation have been met (see chapter 4.7.1), G Data Security Client will automatically be deployed.

7.3. Signature and program file updates

The signature-based protection layers of the G Data security solution need to be supplied with updated virus signatures in order to recognize the latest threats. Each client can independently initiate the update procedure, checking the server for new signature files, or virus signatures can be pushed to the clients. Depending on the wishes of the network administrator, clients can check the network's central G Data ManagementServer, or directly contact G Data's hosted update servers. For enterprise networks, it is recommended to configure the clients to obtain signature updates from G Data ManagementServer. This allows for more control in case an update needs to be rolled back and reduces the amount of network traffic. In addition to signature files, clients can also receive updated program files. G Data Security Client regularly receives updates to improve protection. Unlike virus signatures, program files cannot be obtained from G Data's hosted update servers. To reduce server load, virus signatures and program files update can be distributed using a peer-to-peer system, enabling outdated clients to obtain files from already up-to-date clients.

Two steps need to be configured to enable automatic virus signature and program file updates for all clients. First, ManagementServer should obtain update files from the G Data update servers regularly. Optimally, this happens automatically – if that is not possible, updates can be carried out by connecting to the G Data update server manually or by performing an offline update. Second, the updated files should be distributed to the clients. This too should be an automated process, with several settings available to optimize distribution and prevent network load spikes.

7.3.1. Obtain updates

ManagementServer needs to obtain the latest files from G Data's update servers. This action can be manually initialized or scheduled to happen regularly using the INTERNET UPDATE window of G Data Administrator. The recommended setting is to check the update servers for new virus signatures every hour. Program file updates are not released on an hourly basis, so the update interval can be set to a daily or weekly check. Even for a server that is not connected to a permanent internet connection, scheduling virus signature and program file updates is recommended. The option ON INTERNET CONNECTION carries out an update only when G Data ManagementServer detects that the server has an internet connection available. There is no reason not to let G Data ManagementServer automatically update signatures and program files, other than cases where an internet connection is not always available. To prevent files from being distributed automatically to clients, the Client settings module is the most appropriate place.

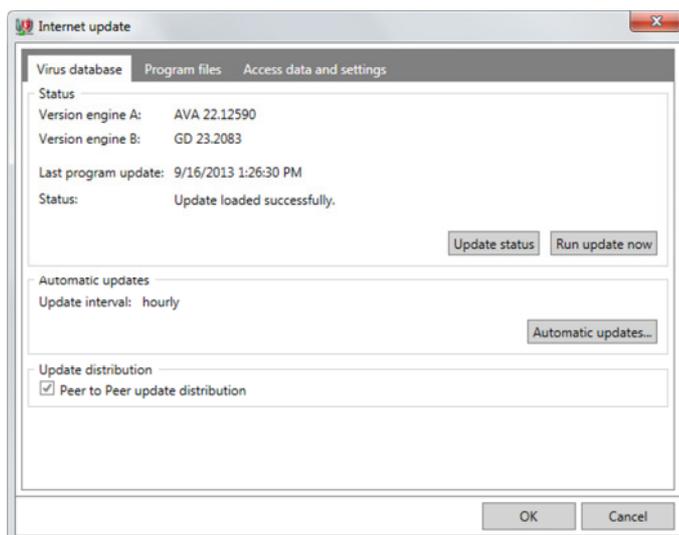


Image 33: G Data Administrator, Options, Internet update, Virus database

For some ManagementServers, a permanent internet connection might not be available for practical or security reasons. While a ManagementServer cannot automatically update itself without an active internet connection, it can still receive the latest server program files, client program files and virus signatures by using offline updates. The Internet Update tool (see chapter 4.5) can be used to apply update files that have been copied from another ManagementServer to the server without an internet connection. This way, updates can be applied by copying them onto a USB stick or burning them to a CD or DVD. Ensure that a ManagementServer with internet connection has the latest updates, either by using the Internet Update tool or via G Data Administrator. Navigate to the folder C:\Users\All Users\G DATA\AntiVirus ManagementServer\Updates and copy the following subfolders onto a USB stick or burn them to a CD or DVD: BB, bd, Client, exchg, GD_SIG, ScanWL, and Server. On the target ManagementServer, insert the CD, DVD or USB stick, and copy the update folders into a

temporary folder (for example, C:\Updates). Open the Internet Update tool and check the option OFFLINE UPDATE. This changes the behavior of the three update buttons: when starting a VIRUS DATABASE, PROGRAM FILES (CLIENT) or PROGRAM FILES (SERVER) update, a popup will prompt for a folder. Select the folder that contains all update folders (for example, C:\Updates), and click OK. Start with the server program files, to ensure that the server itself is at the latest version, and then update the client program files and virus database.

7.3.2. Deploy updates

The second step is the actual distribution of virus signatures and program file updates to the clients. The GENERAL tab of the CLIENTS SETTINGS module allows administrators to configure update settings for one or more clients (by selecting a group of clients or the ManagementServer as a whole). For optimal security, it is recommended to enable automatic updates of virus signatures. Clients need to have the latest virus signatures in order to recognize new threats. The UPDATE SETTINGS window has various options for virus signature downloads. Clients can connect to the central ManagementServer and obtain virus signatures there. They will check the ManagementServer for updated virus signatures in the interval defined under SERVER SETTINGS > SYNCHRONIZATION. To enable central management (with possibilities such as signature rollbacks), it is recommended to have clients download their virus signatures from the ManagementServer. Alternatively, they can download virus signatures from the G Data update servers, in every case (DOWNLOAD INTERNET UPDATE YOURSELF) or only if there is no connection to the ManagementServer (DOWNLOAD INTERNET UPDATES YOURSELF IF VIRUS SIGNATURES ARE OUTDATED AND NO CONNECTION TO THE MANAGEMENTSERVER CAN BE ESTABLISHED). Clients that do not often connect to the ManagementServer, such as laptops that are not always connected to the company network, should be configured to use the G Data update servers as a fallback possibility. To allow clients to connect to the G Data update servers by themselves, login info needs to be defined in the SETTINGS AND SCHEDULING window. Clients can use ManagementServer's access data, or their own (if available). In this case, the SIGNATURE UPDATE SCHEDULE tab should be used to schedule virus signature updates. Hourly updates are the recommended setting for virus signatures, but for clients that are not always online, the option ON INTERNET CONNECTION can be used.

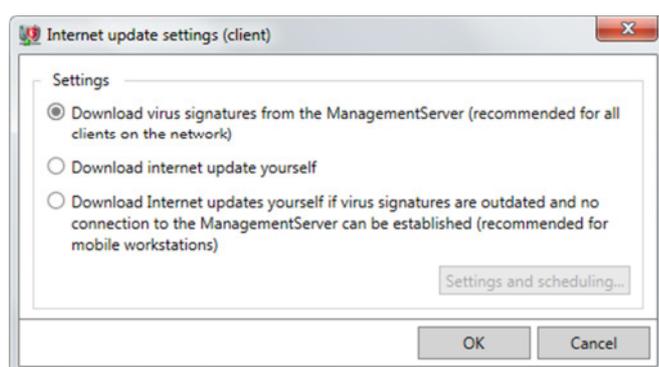


Image 34: G Data Administrator, Client settings, General, Update settings

Program file updates can be installed automatically as well, which helps clients make use of the latest security features of the client software. However, there are more caveats here than with signature updates. An update of the client software may require additional testing to guarantee compatibility with all client configurations in the network. While minor version changes are usually non-intrusive, a staged rollout is recommended. The SOFTWARE UPDATES tab of the SERVER SETTINGS can be used to enable STAGED DISTRIBUTION. This initiates a mathematical calculation where all active network clients are divided into groups (stages). Only after the program file update has been successfully

rolled out to clients in one stage, will the clients in the next stage be served. If there are too many clients where the installation fails, distribution is halted automatically. The number of stages can be defined manually. The larger the network, the more stages should be used to ensure a problem-free deployment. Additionally, the number of days after which the next stage should be deployed can be configured. The default of 3 days allows administrators to check clients for issues, and to halt distribution of a specific update if serious problems are occurring. If required, staged distribution settings can be optimized by editing its settings in the configuration file `Gdmms.exe.config` (see chapter 17.1.6).

Program file updates sometimes require the client to be rebooted. For some client roles, this needs careful planning to make sure machines are not being rebooted in the middle of an important task. The `REBOOT AFTER UPDATE` setting controls client behavior in these cases. The end user can be notified that the client needs to be rebooted, a reboot can be forced or a report can be created in the Reports section, allowing the administrator to manually intervene and reboot the machine at a later point in time.

Distributing updates requires sufficient network bandwidth to be available. ManagementServer and its subnet servers regularly send signature file updates and program file updates to all clients in the network. Update files are incremental and therefore relatively small in size. Still, for larger networks this can cause load spikes. To prevent this from happening, clients can be configured to distribute updates using a peer-to-peer system. This option can be enabled under `OPTIONS > INTERNET UPDATE`. When enabled, ManagementServer distributes updated signature files to a number of clients, which will in turn update each other. With every update that is distributed by the ManagementServer, clients receive information about clients in their vicinity that have not yet been updated and distribute the updates there. The peer-to-peer system can be used without any configuration, but settings can be optimized for specific network situations (see chapter 17.1.5).

Even without automatic updates, clients can still be updated. Administrators can allow the end user to initiate signature updates by themselves (see chapter 8.3). This can be an option for end users that only rarely connect to the enterprise network but want to stay in control of the update process. Alternatively, the administrator can use the Clients module of G Data Administrator to keep tabs on the version numbers and updates that have been deployed to each client. The Clients module shows the version of G Data Security Client, information about the latest updates for engines A and B, and the time at which the client last connected to the server. By sorting and/or grouping the list of clients by these properties, outdated clients can be spotted quickly and updated manually by right-clicking on the client and choosing `UPDATE VIRUS SIGNATURES NOW` or `UPDATE PROGRAM FILES NOW`. An alternative for using the Clients module is obtaining version information by looking at the Dashboard or configuring e-mail reports (see chapter 6).

7.3.3. Rollbacks

Occasionally, virus signature updates can cause problems for specific clients. A generic virus signature may falsely recognize a file as being malicious, or a signature file gets corrupted. The first thing to do is block the specific set of virus signatures (engine update) that is causing the problem. Using the Update rollback function, the update can be blocked. It will no longer be distributed by the ManagementServer, and all clients that contact the ManagementServer will be informed of the block. If a rollback does not take care of the problem, for example when a file got corrupted, a complete update of the virus signatures may be necessary. Normally, G Data ManagementServer and G Data Security Client only download partial (incremental) virus signature updates to reduce

network load and traffic. In the INTERNET UPDATE window, the tab ACCESS DATA AND SETTINGS offers the option VERSION CHECK. By default, it is enabled, causing G Data ManagementServer to compare updates by version number only. By disabling the option and forcing a virus database and/or program file update on the respective tabs, ManagementServer will check the integrity of all update files and re-download them if necessary. The CLIENT SETTINGS module can be used to subsequently disable the VERSION CHECK for the affected client(s), by opening the UPDATE SETTINGS on the GENERAL tab. Initiating the update procedure for that client then supplies a full set of signatures and/or program files.

7.4. Removing a client

If a client is exhibiting compatibility problems, or if it no longer needs to be managed by G Data security software, it can be removed. This process consists of two stages: removing G Data Security Client from the client, and removing the client from G Data ManagementServer's client list. The first step can be carried out by opening the CLIENTS module of G Data Administrator, selecting the appropriate client, and choosing UNINSTALL G DATA SECURITY CLIENT from the CLIENTS menu. This will prompt the administrator to choose if only the client should be removed, or also its associated jobs, reports, messages and backup archives. If a client is removed for compatibility reasons, but should be restored later, it is recommended to keep the associated data. For clients that were installed on hardware that will be phased out, all data can be removed. As an alternative to remote uninstallation, the procedure can also be initiated locally on the client (see chapter 17.2.4).

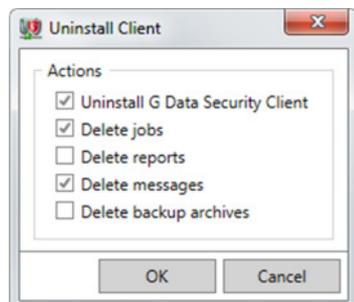


Image 35: G Data Administrator, Clients, Uninstall G Data Security Client

After G Data Security Client has been removed, the client machine will still be listed in G Data Administrator's client management area. It can still be managed, although many functions will only become available after reinstalling Security Client. To completely remove a client from the server, right-click it in the client management area and choose DELETE. After confirming the action, the client will be removed from the list and will have to be added manually if it is to be managed again.

7.5. Performance

G Data ManagementServer can manage a large number of clients. Depending on the size of the network, the number of jobs configured and synchronization settings, performance loss can occur when running a large number of concurrent actions. Several measures are available to resolve performance issues. It is recommended that these measures are only taken when there is a noticeable performance loss. Proactively limiting concurrent client activity is not required if performance has not been impacted.

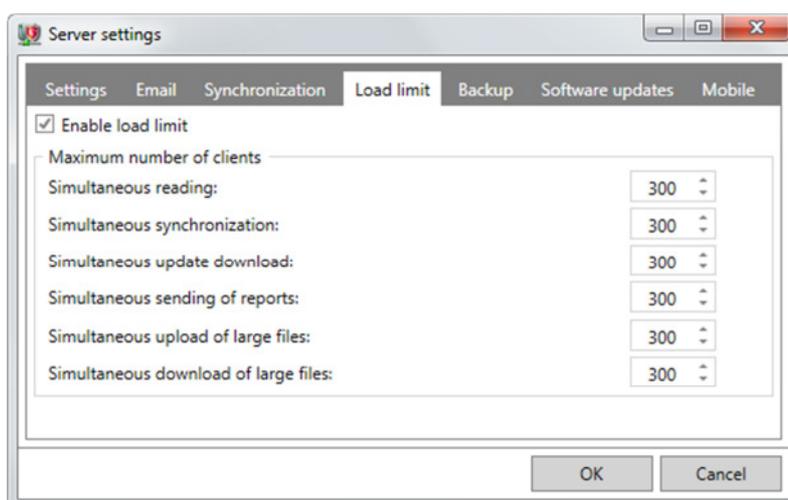


Image 36: G Data Administrator, Server settings, Load limit

The Server settings window offers a load limit option. Different types of actions can be restricted to be carried out for only a maximum number of clients at the same time. If there appears to be a performance loss due to too many client connections at the same time, the load limit can cap the number of connections. On the server side, a large number of update downloads might cause a performance hit if hard disk access slows down the process. Client synchronization can cause higher CPU load on the server, while large file downloads and uploads between clients and server may also cause network congestion. By limiting the number of clients that can concurrently perform these actions, load problems can be mitigated.

For larger networks, it may become necessary to deploy a subnet server (see chapter 3.1). This server installation supports the main ManagementServer. Clients can be allocated to a subnet server and will connect to that server to obtain virus signatures, reducing the server load on the main ManagementServer and the network traffic between the clients and the main ManagementServer. Especially for local branch offices, running a subnet server is recommended: to keep server-client traffic at optimal speeds, the dependency on contact to a main ManagementServer over a WAN can be eliminated.

If scheduled tasks or jobs are causing performance loss on the server, try to reduce the number of tasks scheduled at the same time. Especially if clients are set to regularly report their status to the server, performance can be negatively affected. Furthermore, synchronization of settings between server, subnet server(s) and clients can be set to occur less frequently, although that will not significantly boost performance in most cases.

8. Real time protection

Client security is a multi-layered solution, making sure a malware infection can be prevented at several stages. Before malware can infect a PC, it needs to breach network security (such as a firewall) as well as local client security (such as a file system monitor or an HTTP traffic filter). The local security components play a very important role in preventing malware from infecting systems. Unauthorized network connections are blocked by the local firewall (see chapter 13 for more information). Traffic that is allowed through, will reach the client, which will then filter it using several security modules. Real time security modules like a file system monitor, heuristic scanning, behavior-based scanning, and a browser plugin make sure that malware is stopped before it can execute, regardless of attack vector. When infectious e-mails are sent, malware attachments are removed or quarantined directly, without the file ever reaching the file system. A similar approach is taken by internet and instant messaging security, which filters HTTP and IM traffic and blocks it in case of suspicion. If malware reaches the system through other methods, such as removable media, the file system monitor will block it as soon as the malicious file is written to the hard disk or run. Together, the security measures ensure that malware will not be able to infect the system. That means that the measures should be configured together instead of as separate entities.

Although it can be tempting to enable maximum security for all clients across the network, not all security settings can be identically applied to all clients. Some clients have to be configured with maximum security in mind, for others performance is more important. Security options can be configured for single clients, or for groups, depending on which item is selected in the client management area. It is recommended to organize all network clients into appropriate groups, to be able to apply relevant security settings to multiple machines at once (see chapter 7 for more information on managing clients).

Real time client protection should be configured before deploying any clients. Using G Data Administrator, client devices in the network can be located, organized and configured before G Data Security Client is deployed. For each client or group, make sure to configure optimal protection to be enabled as soon as the client is deployed. Ideally, these settings have been optimized for each client. However, this can be difficult without being able to measure the actual performance at this stage. Administrators can choose to deploy a restrictive set of security settings by default, moderating them afterwards if required. Alternatively, baseline security settings can be used to ensure that performance and usability do not suffer before the administrator has had a chance to adjust protection settings following deployment.

Which settings should be applied to which client depends on the software and hardware that are in use on the client, as well as its role (the context in which the machine is used). For most security settings, only a general recommendation can be given. Often, company policies will dictate certain levels of security, forming a useful guideline for the implementation of a security solution. Other times, a case-by-case policy has to be developed per network zone or even on the client level. This can be a process of trial-and-error – supporting the recommendation that a security solution should be gradually rolled out and its settings carefully monitored.

The core security settings for real time protection can be set using G Data Administrator's CLIENT SETTINGS module. The settings can be adapted for each client or group and their intended use. To help configure each client appropriately, the GENERAL tab features the NOTE field. Administrators can add a description of the client or its deployment, or any other note that helps distinguish between specific security configurations.

8.1. Traffic scans

After the firewall, the next line of defense is checking traffic that has been allowed through the firewall. Various types of traffic can be analyzed by G Data Security Client. Instant messaging applications, which often allow users to exchange files, can be checked by enabling the option `PROCESS IM TRAFFIC`. The client can check incoming IM traffic (default: 5190, the port used by the ICQ protocol, including ICQ and AIM). Alternatively, integration can be enabled for Microsoft Messenger 4.7 or later and Trillian 3.0 or later to automatically check incoming file transfers for malware. General web traffic can also be scanned. By entering a port under `INTERNET TRAFFIC (HTTP)`, G Data Security Client scans incoming HTTP traffic (default: 80 and 443/SSL). Browsing websites that contain malicious software, downloading malware or visiting a phishing website will trigger a warning and access will be blocked. Scanning larger files for malware will take a while; to prevent browsers from timing out during the scan, the `AVOID BROWSER TIMEOUT` checkbox can be ticked. It is recommended to leave this option enabled. To prevent delays, a maximum file size can be defined. Any HTTP traffic containing files that are larger will not be scanned. This option can increase performance, but administrators need to make sure that downloads are then scanned by one of the other security layers (such as the file system monitor). If traffic from a certain website does not need to be scanned, such as a corporate intranet, that website can be added to the `GLOBAL WHITELIST FOR WEB PROTECTION`.

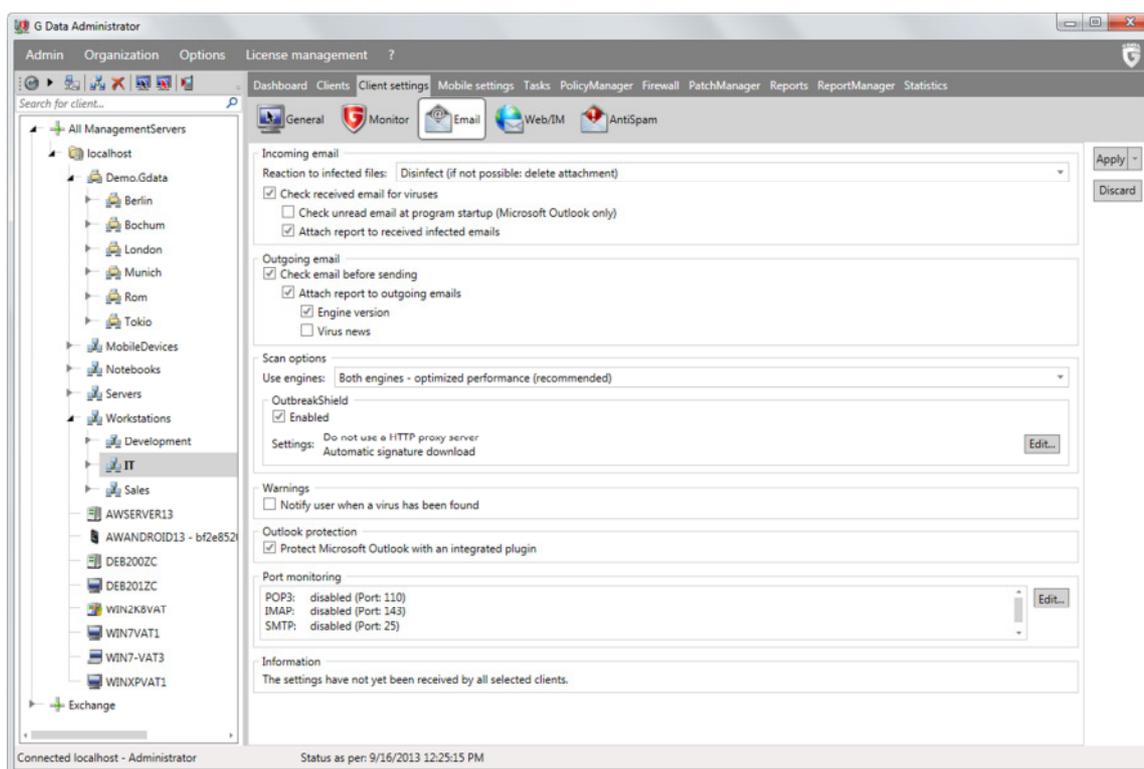


Image 37: G Data Administrator, Clients settings, Email

The third and final traffic category that can be scanned is e-mail (default ports: 110/POP3, 143/IMAP, 25/SMTP)⁷. The ports can be adjusted at the bottom of the settings panel, additionally allowing administrators to prevent e-mail client timeouts by enabling the option for the appropriate ports. Similarly to HTTP traffic, all incoming e-mail is automatically checked for viruses. If malware is found, e-mail messages can be disinfected or have the infected contents removed and a notice about the infection can be appended to the e-mail body. It is recommended to disinfect e-mail messages and

⁷ Note that this only covers client-side scanning of e-mail messages. Server-side e-mail scans can be carried out using G Data MailSecurity.

have infected content automatically deleted if disinfection fails. A more no-nonsense policy is to automatically remove infected content, optionally notifying the user. The LOG ONLY/INSERT WARNING option can be used in networks, network zones or departments where end users need access to incoming e-mail even if they contain infected content or attachments. In that case, malware can still be countered by the file system scanner. However, this is not a recommended scenario, as disabling a layer of protection increases the risk of a malware infection. Networks that use Microsoft Outlook in combination with Microsoft Exchange Server cannot make use of the default port monitoring options, because communication between Exchange Server and clients takes place on several other ports. To ensure that mail traffic from and to the Exchange server is scanned, administrators should enable the Outlook plugin or install MailSecurity's Exchange plugin. In addition to e-mail security scans, G Data can also combat spam. The AntiSpam function checks e-mail traffic and marks it as (suspected) spam. In combination with filter rules in the local e-mail client, AntiSpam helps sort out spam before it even reaches the inbox. AntiSpam does not have a big impact on system performance, but can be disabled if other spam measures are in place, such as a mail server-based solution, or e-mail client spam filters.

Outgoing e-mail should be scanned for viruses to avoid unknowingly sending out malware. Optionally, a report can be attached to outgoing e-mails, indicating that the message has been scanned. This provides recipients with a confirmation that the e-mail message is indeed safe. Under Scan options, administrators can define which scan engines should be used. G Data uses two scan engines for optimal security, but in cases where this affects performance, it can be effective to enable only one of the two engines (see chapter 8.4). As an additional recognition measure, OutbreakShield can be enabled. It uses its own signatures to improve the detection rate of spam campaigns that are distributing malware. By basing itself on general characteristics of the mass distribution of spam messages, detection is already possible even if the virus signatures of the traditional engines have not been updated yet. OutbreakShield is relatively light-weight and should not influence client performance.

8.2. File system

File system protection can be seen as the final line of defense against malware. If a piece of malware reaches the PC and is not picked up on by IM, web or e-mail protection, file system protection will recognize and block it. File system protection in itself is layered, providing protection against known and unknown threats. Whenever files are written to or read from the local file system, the file system monitor kicks in. First, the monitor uses one or two scan engines to scan the file and compare it against the locally saved virus signatures. Known malware will be recognized, at which point the monitor will execute the configured action (such as disinfecting, quarantining or removing the file). If a file is not recognized as malware after comparing it to virus signatures, it will be scanned using heuristic technology. Similar to virus signatures, this approach checks the malware file for patterns commonly exhibited by malware. While heuristics can cause a slightly higher false positive rate, it helps spot malware for which there are no available signatures. Finally, behavior monitoring takes the heuristic approach one step further. While files are being executed, it tracks every action. If the file exhibits behavior that is common to malware, such as excessive write access to the registry or creating auto-start entries, it can be blocked from executing. Additionally, it can be moved to quarantine. In cases where the behavior monitor is triggering false positives, it can be disabled or configured to only log the event. A more specific form of behavior monitoring is BankGuard (available on the WEB/IM tab). It monitors browser system files for Microsoft Internet Explorer, Mozilla Firefox and Google Chrome, and protects against malware that tries to manipulate internet banking

websites.

The different layers of the file system protection can be configured individually. The MONITOR tab provides access to settings for the file system monitor. By default, it is enabled with a balanced profile, providing security without sacrificing too much performance. Administrators can choose to add more security measures, or to increase performance by disabling some. As with most settings, it is recommended to go for optimal security in all cases, unless performance is severely impacted. The MONITOR STATUS should always be set to ENABLED. Completely disabling the monitor deactivates one of the central components of the G Data security solution and should only be considered if an individual client is experiencing extreme compatibility problems, or if it is sufficiently protected by alternative security measures.

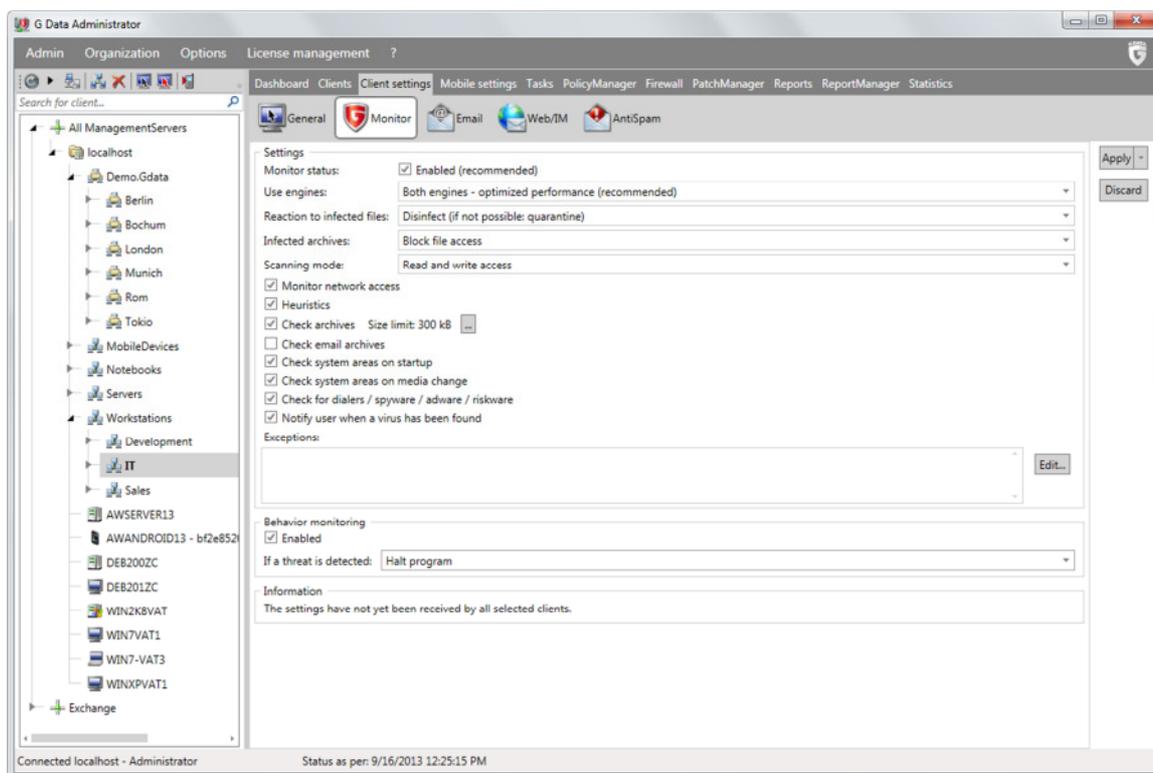


Image 38: G Data Administrator, Clients settings, Monitor

The file system monitor uses two software engines to scan files for malware. This offers optimal security, by ensuring that a malware infection that is missed by one engine will be picked up by the other. The default and recommended option is to enable both engines in optimized performance mode. For optimal security, both engines can be enabled while ensuring that files are always checked by both engines, but this will cause a performance hit. If performance is affected even when using the optimized performance mode, clients can be configured to scan files with only one of the two engines.

G Data can take several types of action when a malware-infected file is encountered. The most restrictive setting immediately deletes infected files. This makes sure that malware never has a chance on the system, but causes data loss if important documents are infected or falsely flagged as malware. It is recommended to choose disinfection as the primary action. This will let the file system monitor try to repair infected files, in order to restore access to it. If that is not possible, the file can be deleted or quarantined or access can be blocked. While blocking access does prevent the malware from running, it leaves the file untouched. Quarantining the file makes sure it cannot be

accidentally run by anyone, while keeping open the option of restoring it or manually repairing it. From the quarantine, administrators can choose to send their file to G Data for further analysis. The file system monitor can scan the contents of archives (such as ZIP files), but does not remove single infected files from an archive. Archives that contain an infected file can be moved to quarantine as a whole or deleted, or access can be blocked. Because archived files will be scanned once they are extracted or opened, there is a reduced risk when working with infected archives. Instead of deleting the file, blocking access should do. For performance reasons, archive checks can even be disabled altogether.

The scanning mode decides when the file system monitor will do its work. Files can be scanned on read access, on read and write access, or on execution. On execution will provide the baseline protection of scanning files when they are run. This will prevent malware from infecting the client PC but, like blocking access, does not help against further spreading the malware. To make sure that infected files are detected even when they are not being run, choose read or read and write access. That way, files that are being copied to or from other folders, disks or clients will also be scanned. This way, malware that is actively spreading itself across the network will be detected when it tries to infect other PCs by writing files to network shares. Scanning for malware on read and write access is a hard disk intensive process. If client performance is impacted, the scanning mode can be scaled down to read access. On execution is only recommended for cases where even read access scans are too system-intensive.

By default, the file system monitor checks archives, e-mail archives and system areas for malware. Due to their file size, archives can be troublesome to scan. A size limit can be defined to prevent the monitor from reducing performance by initiating lengthy on-access archive scans, or archive scans can be disabled altogether. Even when automatic on-access scans for archives are disabled, users will still be able to initiate a manual scan when they encounter a suspicious archive file. When enabling e-mail archive scans, keep in mind that e-mail software will complain if its data file is moved into quarantine. In most cases, it is better to disable e-mail archive scans and let the file system monitor check files once they are extracted from the archive (for example, when saving an attachment to the local hard disk). System areas (boot sectors) can be scanned on start up or on media change. This scan should be enabled in order to detect boot sector viruses. Also enabled by default is the option to search for dialers, spyware, adware and riskware. While not strictly malware, these types of files are generally unwanted.

When the file system monitor detects a threat, it automatically takes the action that has been defined by the administrator. In addition, G Data Security Client can display a message on the client PC that informs the user that malware has been found. The message contains the file, its path and the name of the malware found. Displaying a warning can help the user realize that the current program that is being executed or website that is being visited is malicious. However, some users may be confused by the message. Messages about threats that have already been blocked are not relevant for all users and can be disabled on an individual basis.

Specific files or folders can be excluded from the file system monitor scans. Using exceptions, files that are impractical for on-access scanned can be ignored. For example, large files that are seldom used do not need to be scanned on-access if they are included in a scheduled scan job (see chapter 9.2.1.3). Similarly, database files can be excluded from the on-access scans if they are regularly scanned by an on-demand scan. It is recommended that not too many exceptions are added. Only files that impede performance when scanned with on-access scans should be defined as an exception, and only if they are regularly scanned by a regular scan job. Files that generate false positives but are known to be safe can be added to the exception list as well, after it has been

verified that they are indeed not malicious. When defining exceptions, extra care should be taken that they only apply to those clients or groups that really need the exception. Since an exception prevents files from being scanned for malware by the file system monitor, it should be limited to as few clients as possible. The exception list can be populated with directory, drive, file and process exceptions. Directories, files and processes can be manually entered into the text field or selected from a folder tree. Selections can be made from local folders and files or by opening the folder tree on any of the clients in the network. For files and processes, the full path and file name should be entered in the text field. Exceptions can be defined using the wildcard symbols ? and *, representing single characters and entire character strings, respectively.

The other layers and components of file system protection can be enabled and disabled separately. Network access monitoring is enabled by default, but can be disabled if the computer is located in a network where all clients are being protected by G Data. Heuristics and behavior monitoring are essential security measures, and should only be disabled if they are generating too many false positives.

8.3. End user security permissions

Security settings are managed centrally by the administrator. However, in some cases it can be useful to let end users initiate virus scans or change settings. Other times, administrators may need to access local security settings while responding to a support call. The GENERAL tab of the CLIENT SETTINGS module lets administrators grant several permissions to be changed locally using the system tray icon of G Data Security Client. These options can be password-protected, a practical possibility if not all end-users on the machine should be granted access, or if only the administrator should be allowed to change settings.

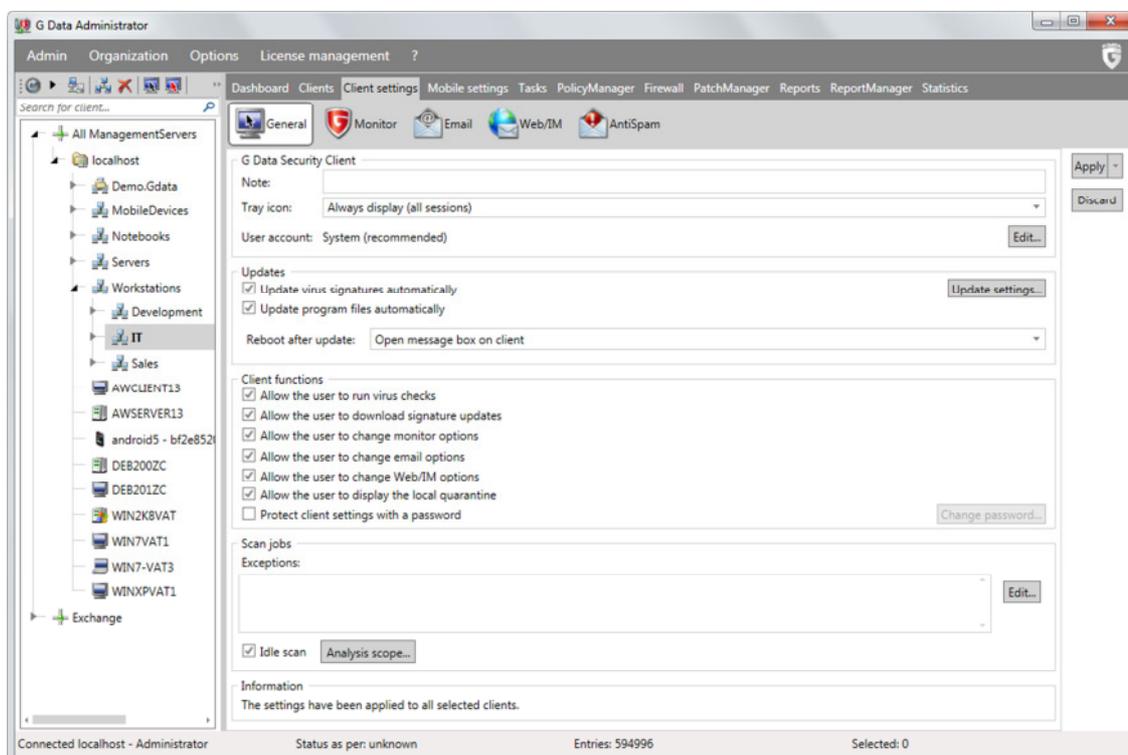


Image 39: G Data Administrator, Clients settings, General

To reduce the number of support calls in case of a possible malware infection, end users can be

allowed to check files or folders for malware manually. By selecting the option `ALLOW THE USER TO RUN VIRUS CHECKS`, a manual virus scan can be initiated, regardless of the server-side scan schedule. Individual files or folders can be scanned, as well as the computer as a whole, or removable devices, or the memory and startup areas. Clients and end users that are expected to work with documents from unverified sources (such as internet downloads or files from removable devices) can use manual scans for an extra verification of file safety, but in most cases, the centrally defined monitor settings should already scan the files and take care of any possible infection. If the option `ALLOW THE USER TO DOWNLOAD VIRUS SIGNATURE UPDATES` has been enabled, end users are allowed to update the local virus signatures independently from the update schedule. This can be useful for end users that regularly use their devices while not being connected to the corporate network, such as laptops. On the other hand, it can lead to fragmentation in the number of signature file versions available across the network, complicating client management and troubleshooting.

All security options for the file system monitor, e-mail scan and Web/IM scan can be made accessible to the end user. This possibility should be handled with care. In order to keep a unified security policy, it is recommended not to let end users change security settings. All changes should be authorized by the administrator and tested beforehand. By enabling one of the appropriate options, all security measures can be disabled. Only if administrative access to the settings is required locally when servicing a client, the option should be enabled, and only if a password is set.

G Data Security Client can offer access to the local quarantine. If the administrator has defined rules for infected files to be moved to the quarantine, end users can have a look at the files that have been isolated. For each infected file, the date and time of detection are listed, as well as the full file and folder path, and the name of the virus. Files can be manually disinfected, deleted or moved back. This last option is the reason that access to the quarantine should only be granted to experienced users or administrators. Files that are moved back without disinfecting them first can pose a risk to the system.

When using the PatchManager module, administrators can choose to allow end users to view patches that have been installed on the system or patches that are available but have not been installed yet. This option does not pose a security risk, and can be useful if end users encounter problems after a patch installation and want to file a rollback request, or if a patch should be deployed with priority to fix an end user's compatibility problem. More information can be found in chapter 15. Administrators that have enabled G Data Firewall can allow users to enable or disable the firewall, and to change the off-site configuration. Letting users disable the firewall at will is not recommended, unless there is a strict security solution in place on the network level. Changing the off-site configuration is useful for clients that are often connected to networks other than the enterprise network. Users will be able to create rules and rule sets for those networks, which will automatically be overruled when connecting to the enterprise network. More information about rule sets can be found in chapter 12.

Note that permission to change G Data Security Client settings is not related to application permissions, device permissions, web content control or internet usage time. Permissions for these types of actions can be granted or revoked using the PolicyManager module (see chapter 14).

8.4. Performance

Client hardware and software, as well as network infrastructure, provide only finite capacity. Traditionally, information security has always had to be established at the expense of performance. The more security measures that had to be implemented, the more performance impact could be

measured. For most modern client-server security solutions, therefore, performance is a key factor. G Data offers optimized security modules that hardly affect client performance, even when configured for optimal security. Different circumstances, such as hardware specifications or software deployments, can require a different balance between security and speed. Having built a network according to defined network zones and client roles (see chapter 1.1) helps decide which clients need more security and which clients can focus on performance. Finding the correct balance is difficult: for some enterprise networks, as many security layers as physically possible should be enabled; for others, client speed is more important. The advantage of the multi-layered approach that G Data offers is that the security functions can be optimized for each client. Several options complement each other, allowing one or two to be disabled without affecting performance.

In general, it is recommended to err on the side of caution and think twice before disabling a security feature for performance reasons. When deploying a security solution, start with the maximum level of security that is appropriate for that client, and reduce the security level only if the client is noticeably impacted. Some of the security settings impact performance more than others. The file system monitor can be demanding if the client is mainly used for file operations. Settings like archive scanning can delay the processing of large files, while read and write access scanning can cause significant slowdowns for mechanical hard drives. Clients with a low-end CPU can experience lag when using both scan engines – disabling one of them will significantly increase performance. Other file monitor settings, such as heuristics, or behavior monitoring, only minimally impact client performance.

Usability is impacted by other factors than performance alone. Security software always has a small risk of identifying regular files as malware, based on patterns in the heuristics module or behavior monitoring. If protection has been configured too tightly, more false positives will emerge, but if security is lax, malware may slip through unnoticed. As with performance problems, false positives should only be remedied if they are actually occurring. If a file is identified as malware, the `REPORTS` section of G Data Administrator will show on which client it was detected and by which protection module. Administrators can choose to change settings of the relevant module to be generally less strict, or whitelist the specific file (if the security module supports this measure). Because generally lowering security might put clients at risk, at first the affected file should be whitelisted only. This prevents it from being detected as malware and removed or quarantined. If a specific security module often wrongly identifies malicious files, it can be configured for lower security. Completely turning off a module should only be a last resort. In most cases, quarantining a file instead of deleting it, or logging an infection instead of directly acting upon it will do. However, these cases do require administrative attention. If G Data is configured not to take care of an infection automatically, manual investigation will be necessary. Alerts can help notifying administrators of such cases (see chapter 6.2). If changing security module settings is not possible, for example when a specific group of clients requires absolute security, PolicyManager can help restricting access to applications, devices or websites (see chapter 14).

8.5. Operating system security

Especially the more recent Windows client operating systems are based on a solid security architecture, starting with Windows Vista. While security software significantly reduces the chance of malware infections, additional measures can be configured on the operating system level. Some aspects of both network security and local operating system security can be efficiently controlled using local settings or network-based group policies. Administrators can choose to enforce these settings on all clients, or to protect only the most vulnerable network zones and client roles.

By default, Windows is configured to run many background services, essential and non-essential. Some are supporting important Windows functions, others are only used for very specific configurations or software suites. Every service, however, can be an attack vector. Attackers are constantly looking for vulnerabilities in Windows services to exploit a security hole and gain access. Patching the affected services helps, but can only be done as soon as Microsoft or the software vendor locates the security hole and issues a patch. Unused services can be disabled in advance, thus taking away the possibility that a vulnerability is abused by hackers. Furthermore, disabling services can lead to performance increases. Running the command *services.msc* on a Windows PC will open the SERVICES manager that displays all services that have been installed on the system, whether enabled and running or not. Disabling services is not a quick win measure; administrators should make sure that the client PCs in question do not depend on a service before disabling it.

To decrease the risk of malware spreading via removable media (such as USB sticks or CD-ROMs), Microsoft has replaced autoplay in recent versions of its operating system with a window that allows the end user to pick a specific action to be taken, such as opening the medium's root folder or importing data from it. This circumvents the often-used autorun.inf configuration file. While this effectively stops malware that spreads itself through autorun.inf configuration, users can override it. In the window that pops up whenever a removable medium is inserted, users can select a default action which will be carried out each time a medium of that type is inserted. If an infected USB stick is inserted into the computer at a later stage, it will be run just like any other medium. With a security solution installed, even if malware is run it will not cause any damage. However, for PCs in corporate networks, disabling autoplay completely will increase security. For smaller networks, using local registry settings to disable autoplay can be practical: using the Registry Editor, add a DWORD value named *NoDriveTypeAutoRun* with hexadecimal value *0xFF* to the key `HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\Policies\Explorer` to disable autoplay for all types of drives. Alternatively, a hotfix is available from Microsoft to change the setting automatically⁸. For corporate networks, a group policy is available to disable autoplay for specific machines or users.

Downloaded files and e-mail attachments are another common source of infections. For this reason, browsers and e-mail clients usually warn end users before downloading any files. This helps educate about the risks of downloading files from public sources, in theory. In practice, the messages are often dismissed without looking. Security solutions pick up on malware long before the file hits the hard drive (through web protection, blacklisting and other solutions), and even if it does, the file system monitor will recognize and block it. Still, some administrators may want to rule out the risk of malware even being pulled in by a client browser. Using Security Zones, file downloads can be blocked across all browsers that support the built-in Windows technology Windows Attachment Manager (Microsoft Internet Explorer, Google Chrome). Administrators can configure URL actions to be disallowed in certain security zones, such as executing a download⁹. For other browsers, security zones may not be honored. In that case, limiting file save access can help, as well as enabling User Account Control or an appropriate group policy.

Other basic security musts include limiting the permission level for end users to the lowest level possible. When an end user does not need advanced system permissions to perform tasks on the client PC, these permissions should not be granted, to prevent them from being abused by malware. This includes not letting any user log on to a client PC using an administrator account, and using group policies to restrict access to local and network resources. Blocks for other actions, such as

⁸ See <http://support.microsoft.com/kb/967715>.

⁹ See <http://msdn.microsoft.com/library/ms537183.aspx>.

viewing removable media, or using certain applications, can easily be configured using G Data's PolicyManager module (see chapter 14).

There is no one-size-fits-all solution where local security settings are concerned. Each network zone will have different demands for its client PCs, and each individual user may need further changes to security policies. A good starting point for networks that are using Active Directory is to explore the possibilities of group policies. This built-in feature allows administrators to configure extensive rules for the use of local and network resources and for managing permissions. For networks that are not using Active Directory, many settings can be adjusted locally.

9. On demand protection

In addition to real time protection, client PCs can also be protected on demand. On demand scans are carried out once or periodically and check a predefined area on the client PC for malware. On demand protection should be configured in addition to real time protection, as the latter type does not scan files that are not being accessed at that time. On demand scans detect malware in all files on the client PC, regardless of whether there is current read or write access. By planning a recurring on demand scan for a client PC, its entire hard disk can be scanned to check if malware has infected otherwise dormant files. Two types of on demand scans can be planned using G Data software: idle scan or a single or periodic scan job. The idle scan is the recommended type of on demand protection, as it does not need to be scheduled: it automatically scans clients whenever they are not in use. Alternatively, scan jobs can be used to scan clients at a predefined moment in time.

9.1. Idle scan

A traditional full scan job, planned at a specific time of the day, requires a significant amount of computer resources, making it impractical to run a scan while an end user is logged on. Scheduling a full scan during off-peak hours prevents productivity interruptions, but requires the client to be powered on. Configuring the idle scan functionality provides an ideal solution: when the client is powered on but not in use, G Data Security Client starts an automatic scan in the background. It will scan predefined drives, files and folders and automatically pause when the user returns to the client. This provides the functionality and level of security of a regular scan but prevents the performance loss that can accompany it.

For clients that are not always connected to the company network, idle scan can be a replacement for scheduled scans. A laptop that connects to the network after having been disconnected for a longer time will receive newly planned scan jobs and immediately execute them if they are overdue, which can lead to reduced performance. Enabling idle scan for clients like these, and disabling the scheduled full scan, will make sure that they receive full protection while avoiding being overburdened at startup.

To determine whether the client is idle, the G Data Security Client tray icon needs to be enabled for all sessions. Idle state is determined by several parameters. The scan will only be started if the end user has not been using the client for at least one minute, and never in the first ten minutes after the client has been powered on. If the user is not at the PC, but background tasks generate CPU or I/O activity, the idle scan will be paused. Furthermore, scheduled jobs will be carried out before the idle scan can commence. If the idle scan is running, but any of the parameters change (e.g. when the user returns to the client, or a scheduled job starts), the idle scan will be paused. When the parameters are met next time, it will pick up where it left off. For the end user, an idle scan will work just like any other scheduled scan. It does its work in the background, and notifies the user if malware is found (if this setting has been enabled by the administrator). For administrators, an idle scan will not generate a regular report about scan start or stop, but any malware will be reported in the REPORTS module as usual. After the idle scan has completed one full run of scanning the specified folders, it will automatically start anew after seven days.

Idle scan can be enabled on the GENERAL tab of the CLIENT SETTINGS module. For some clients, the regular quick and full scans may suffice; for others, an idle scan can provide security while reducing the time and performance required for a regular full scan. As with other settings, it can be enabled or disabled per client. It is recommended to enable idle scan for all network clients, disabling it only if it leads to performance issues or other problems. The ANALYSIS SCOPE can be defined per client and

ranges from all local hard drives to individual folders. For high-risk clients, idle scan can be configured to monitor sensitive folders. Alternatively, it can partly take over the tasks of the full scan, and scan all local hard drives. The scan settings for the idle scan are identical to the Monitor scan settings and will be taken from the MONITOR tab. This concerns the engine settings, REACTION TO INFECTED FILES, INFECTED ARCHIVES, and SCANNING MODE, as well as the scanner specifics MONITOR NETWORK ACCESS, HEURISTICS, ARCHIVES, CHECK E-MAIL ARCHIVES, CHECK SYSTEM AREAS, CHECK FOR DIALERS, and NOTIFY USER WHEN A VIRUS HAS BEEN FOUND. Since the idle scan is a special kind of scan job, the exceptions for the idle scan are taken from the scan job exceptions list on the GENERAL tab.

9.2. Scan jobs

Scheduled on demand protection is carried out in the form of single or periodic scan jobs. Single scan jobs are run one time only, while periodic scan jobs are repeated according to the defined schedule. Both types of jobs can be planned and managed from the TASKS module. As with other modules of G Data Administrator, the jobs that are planned apply to the client or group selected in the client management area. The module shows a list of currently defined jobs. By default, all types of jobs are shown, including backup jobs, patch management jobs, etcetera. To display only scan jobs, click the appropriate button in the module's toolbar. For each job, several properties are listed. For client jobs, the client for which they have been defined is displayed (for groups, the group name). The STATUS column shows the current status of the job. For group jobs, the status can be checked per client by selecting the appropriate client on the left. If a job has been run at least once, the job listing can be expanded to show the associated scan log(s). Double click on the log to view a detailed list of job results. The INTERVAL column displays the defined scan interval, such as Once for a single scan job, or Daily for a periodic scan job that runs every day. Finally, the SCOPE displays the analysis scope that has been defined for the job.

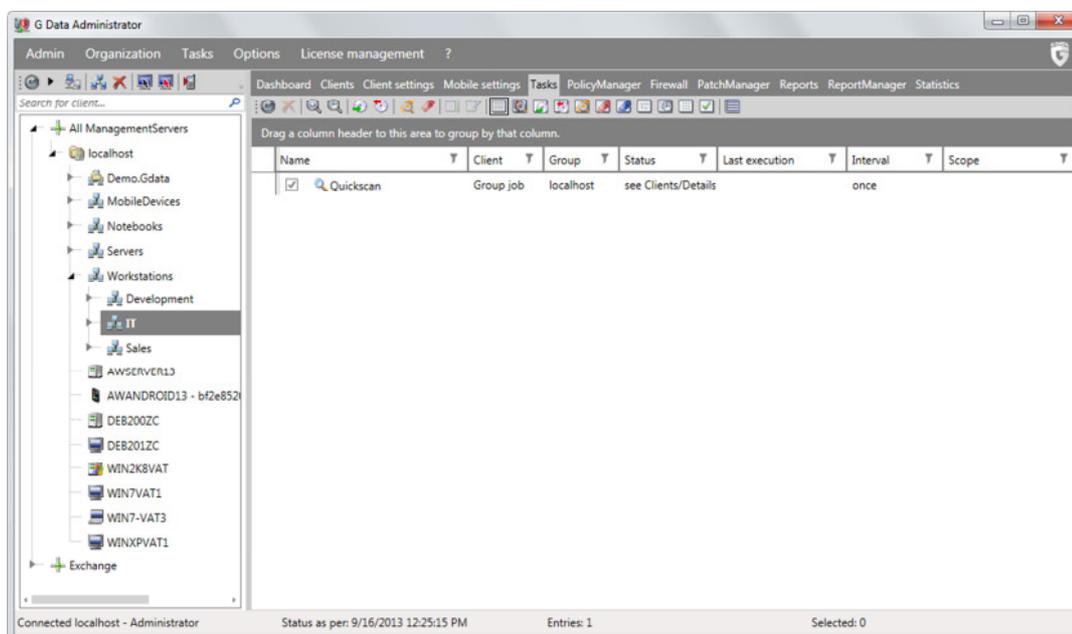


Image 40: G Data Administrator, Tasks module, Scan jobs

Single and periodic scan jobs can be defined by clicking the respective buttons on the task bar, or by opening the TASKS menu and choosing NEW > SINGLE SCAN JOB or PERIODIC SCAN JOB (see chapters 9.2.1 and 9.2.2 for more information about scheduling scan jobs). After defining the job, it will immediately show up in the TASKS list. The STATUS column will show its current status. Once the job has been

executed at least once, the date and time of its last run will be displayed, as well as a scan log (accessible by clicking the plus icon in the leftmost column). If any malware is found during a scheduled scan, the action that has been defined when setting up the job will be automatically carried out. The scan log will show an entry about the infection, and a report will be added to the REPORTS module (see chapter 6.2).

For both single and periodic scan jobs, some general options can be set on the JOB SCHEDULING tab of the scan job window. End users can be allowed to pause or cancel a scan job. In cases where a scheduled scan might interrupt end users' work, pausing a scan will allow them to continue working without a performance hit. This option should be used with care: system security can be compromised if an end user decides to cancel an essential scheduled scan. End users can be notified when a virus is found during a scheduled scan. While this can be a useful option for file system monitor scans (see chapter 8.2), the end user does typically not need to be notified during a scheduled scan. The scanner takes care of the infected file automatically, and the infection will be reported in G Data Administrator's REPORTS module. During the scan, G Data Security Client can report its progress to the ManagementServer every two minutes (updating the details of the scan job in the TASKS module). This can be valuable for the administrator to keep an eye on the exact progress of a scan, but is rarely useful for recurring scans. Only if a one-time scan is executed immediately, keeping track of its findings directly can be necessary. Clients can be shut down automatically after a scan job is completed. However, to prevent data loss or otherwise unexpected behavior, an automatic shutdown is not possible when an end user is logged on to the client when the scan is completed. Automatic shutdown is useful mostly if scans are scheduled after a regular work day or on weekends, when the clients are not scheduled to be used directly after the scan. Periodic scan jobs can be delayed if the client is not powered on at the scheduled time. By selecting this option, it is ensured that no jobs are being skipped. By skipping a job, an infected file may persist on the system until the next scan, possibly being further distributed to other systems. This option should only be disabled for jobs that are run regularly, so that the next job will be carried out relatively soon after the skipped job.

9.2.1. Periodic scan jobs

Before planning any periodic scan jobs, it is recommended to think of the scan schedule as whole. Every network client should be scanned regularly, but planning one single scan job including every PC in the network is not productive. Scheduled scans for servers should be different from those used for clients, and different client roles can require different scan settings. Multiple scan jobs can be planned for the same client: for example, a daily quick scan and a weekly full scan. As a whole, the scan schedule should ensure that none of the clients go without scan for a longer amount of time. As addition to real time protection, regularly executed scan jobs will ensure that client PCs are completely malware-free. However, scan jobs can cause performance hits, depending on the number of files that needs to be scanned. A daily scan of the complete hard drive is therefore not recommended: it may take long and requires CPU power and semi-constant hard drive access. For most client PCs, the idle scan (see chapter 9.1) is the best option. If the idle scan is not being used, a weekly full scan and daily quick scan, combined with real time client protection, will find and remedy all malware infections. For client PCs that have a higher risk of malware infection, for example when it is often used to handle downloads, a full scan can be scheduled more often, if client performance permits. Servers, such as file or database servers, may be under a significant load, leaving hardly any CPU cycles for a malware scan. For that type of machines, it may be necessary to schedule scans during off-peak times or preset maintenance windows.

For almost every enterprise network, several scan jobs should be configured. All of them should recur periodically, optionally supplemented by one-time scans to cover exceptions or acute cases. The following paragraphs will define several types of periodic scan jobs that can serve as templates for the specific needs of an enterprise network. For all periodic scan jobs counts: there is no one size fits all solution. Adapt the settings to the enterprise network, down to network zone and individual client levels, if necessary. After the job has run a few times, check to see if it produces the expected results and does not impact client performance too much.

9.2.1.1. Full scan

The first and arguably most important type of periodic scans is the full scan. In addition to real time protection, each network client should be scanned regularly by a scheduled scan job. A full scan picks up malware that is not being detected by the real time client protection. While a file system monitor detects all malware that is being read, written or executed, it does not proactively scan files that have been saved to the hard drive. If a file has been written to the hard drive while the monitor was not (yet) enabled, the file system monitor will only detect it once the system attempts to open or run it. Because it scans the complete hard drive, a full scan will detect malware before the system tries to execute it. The more often a full scan is carried out, the higher the level of security and certainty that no infected files are located on the system. However, because it scans all files on the system's hard drive(s), a full scan is very performance-intensive. On most systems, running applications will slow down and end user work will be impacted. For performance reasons, a daily full scan in most cases will not be an option. Low- and medium-risk clients can be set to be scanned weekly, after work hours or during the weekend. High-risk clients could be scanned more often, but ideally only when the machine is not in use. Servers should be scheduled for full scans as well – if performance is an issue, a full scan can be planned during a server's regular maintenance window.

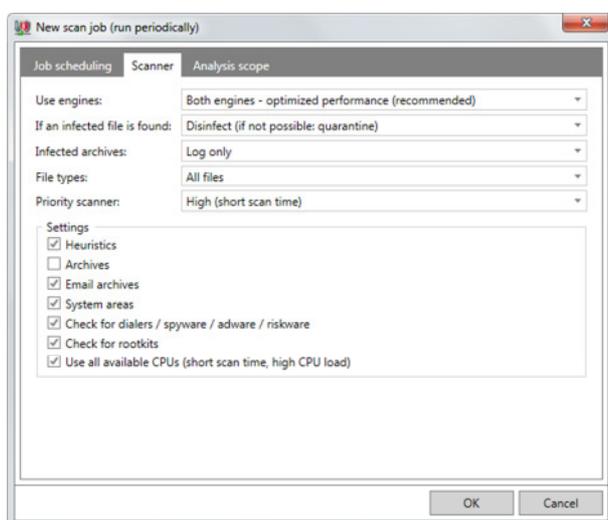


Image 41: G Data Administrator, Tasks module, New periodic scan job (Full scan)

A full scan can be defined as a periodic scan job in G Data Administrator's TASKS module. Make sure that the appropriate client or group for the scan is selected in the client management area. The PERIODIC SCAN JOB window opens on the JOB SCHEDULING tab. Select the day and time at which the scan should be run for the selected client(s). Since a full scan requires resources and can take a considerable amount of time, make sure to plan the job so that it does not overlap with other tasks (periodic or single scan jobs, backups or PatchManager tasks). The SCANNER tab defines the parameters with which the scan job will be carried out. Because a full scan already requires

significant client resources to be available, and should always be run while the client is not in use, the scan settings do not have to be tweaked for performance. However, settings should still be optimized for the client(s) that the job is planned for. The safest scan engine option is to use both scan engines. Even if it impacts client performance during the full scan, using both engines is recommended: it provides optimal malware recognition. Which action is to be taken when an infected file is found depends on the administrator's preferences. Generally, disinfection is recommended, with the file being moved to quarantine if disinfection is not possible. This will prevent false positives from being deleted immediately and allows administrators to further examine files in the quarantine. Infected (e-mail) archives can be moved to the quarantine as well, but require more attention. Even if only one file or e-mail within an archive is infected, the complete archive will be moved. Files that have not been infected but have been archived together with an infected one will be quarantined too. Deleting the complete archive is an even more drastic option, and not recommended. Infected archives can be logged, in which case the administrator should make sure to check the Reports module regularly. All files should be scanned; changing the file types to be scanned to ONLY PROGRAM FILES AND DOCUMENTS will disregard potentially infected files. Scanner priority can be set to High if the client will not be used during the scan; the scan time will be significantly reduced. Otherwise, the Medium or Low settings should be used. The specific checks, such as Heuristics, E-mail archives, System areas, Dialers, and Rootkits, should all be enabled to ensure maximum protection. Archives can be included in the scan, but since they are often relatively immutable, a separate, they could optionally be covered by less frequent periodic scan (see chapter 9.2.1.3).

The analysis scope for a full scan should include all hard drives. By selecting Check the following directories, administrators can select one or more folders that should be scanned, allowing for a more granular control over the folders to be included in the full scan. While this option can be used to exclude folders that should not be scanned, it is recommended to leave Check local hard drives selected and define eventual exceptions using the Client settings module (see chapter 9.3). In addition to the full hard drive scan, a memory and autostart scan should be defined separately (see chapter 9.2.1.3).

9.2.1.2. Quick scan

The second type of recurring scan job is the quick file scan. This scan job should be configured to run daily for all clients in the network. Servers can be included in the job as well, but only if there are enough resources available to run the scan. A quick scan checks the files with the highest risk of infection, skipping files that cannot be executed (thus cannot possibly infect the client). Reducing the number of files to be scanned significantly reduces the amount of time needed for the scan, while retaining a proper level of security. Like the full scan, a quick scan can be configured by choosing PERIODIC SCAN JOB. A quick scan should run daily (weekends can be omitted for clients that are only powered on and in use during the week). The time at which the job should be run can be freely defined. Quick scans are more lightweight than full scans, but can still cause a performance hit (depending on the settings on the SCANNER tab). Planning a scan during lunch time is a possibility, or after the typical office hours (if the clients are not powered down). Depending on the client hardware configuration and quick scan settings, a typical end user may not notice any reduced performance, allowing a scheduled scan time at any point during the day.

The Scanner tab configures the types of malware scans to carry out. For a quick scan job, these settings will mostly depend on the required speed. A quick scan is meant to be a light-weight daily scan that checks executables for malware, and does not burden the client too much. For optimal

detection it is recommended to use both scan engines. If a client's hardware or software configuration causes performance problems during the quick scan, scans can be carried out using only one engine, but this will reduce the detection rate slightly. Like a full scan, the action to be taken if an infected file is found can be set according to the administrator's preferences: disinfect/log is the recommended option. For a quick scan, set the File types to include Only program files and documents. This will scan only those files that can actually infect the system by being executed or opened. Other files will be ignored, saving time. Scanner priority depends on the resources that are available to the scanning process. If the quick scan will be run while an end user is likely to be working at the client, a medium or low priority should be set to allow for normal performance. Using this setting, however, the scan will take longer to complete. For the quickest result, the priority can be set to High, but this will impede performance. Similarly, the Use all available CPUs option can speed up the scan, but might cause a significant performance drop. Enable the Heuristics setting to carry out the scan using pattern technology to recognize malware patterns. For a quick scan, Archives can be ignored by deselecting the checkbox: since they cannot be executed, they do not need to be scanned in a quick scan. The same goes for e-mail archives. Note that in both cases, clients should have the file system monitor enabled, to make sure that malware contained within archives is blocked as soon as it is extracted. System areas should be scanned during a quick scan and Check for rootkits should also be enabled. Both options are not very performance- and time-intensive, yet scan important parts of the client system. Finally, the dialers/spyware/adware/riskware scan can be enabled if they are considered malware by the administrator, or disabled to save (some) time.

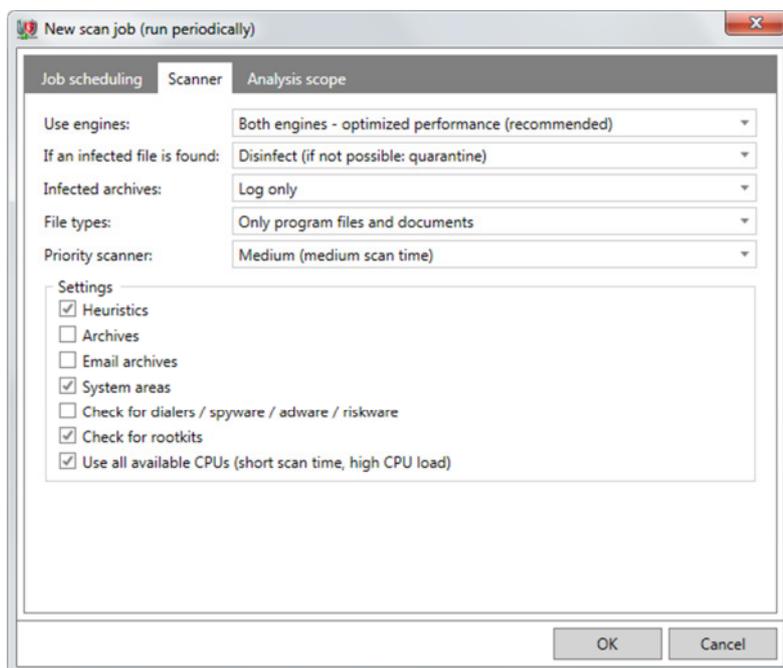


Image 42: G Data Administrator, Tasks module, New periodic scan job (Quick scan)

For a quick scan, the analysis scope should include all client files. Select CHECK LOCAL HARD DRIVES to run the scan job on the complete hard drive. It is possible to limit a quick scan to pre-defined folders, such as the Program Files folder or the Windows system folder. This will reduce the time needed to complete the scan, but represents a security risk, since malware executables can be stored in any folder – in fact, many varieties try to avoid detection by using uncommon folders.

9.2.1.3. Other periodic scans

In addition to the typical full and quick scans, there are several other types of common scans. Most importantly, administrators should plan a memory and autostart scan. This will check the often-targeted Windows autostart areas and the currently running programs in memory. It is a relatively lightweight measure that can be planned for all network clients by defining a new periodic scan job with the ANALYSIS SCOPE set to CHECK MEMORY AND AUTOSTART. Typically, a memory and autostart scan is run on system startup, but it can be planned as an hourly or daily job as well. System startup is a good moment for other light-weight scans. Large scan jobs should be avoided: systems are usually powered on by an end user, and should be ready for work as soon as possible. Delaying startup with a malware scan is often not acceptable. However, if a client has high-risk file or folder set that should be scanned regularly, scheduling a job that will check the files after startup, before they are being used, can provide additional security.

Periodic scans can be used to pick up file types or folders that are not scanned by the regular quick or full scans. Large archive files, for example, are typically large and not altered very often. Even if an archive contains an infected file, the client will still be secure: as long as the file system monitor is enabled, the file in the archive will be blocked once it is extracted. Because of the various protection layers, archive files can be excluded from the regular full scan to save time. A separate periodic scan can be scheduled to scan archives. Whereas the full scan would typically occur weekly, an archive scan could be planned to be carried out only every second week, or once per month. Alternatively, periodic scans can be scheduled in between full or quick scans, to check a specific folder for malware. Especially high-risk clients can profit from an extra scan (or an increased scan frequency for the regular full scan).

9.2.2. Single scan jobs

Scans can be scheduled to occur only once. Unlike periodic scans, these single scan jobs are carried out once, after which a report is saved and the job will not be repeated. Single scans can be a powerful tool for administrators that are locating or cleaning up malware. When acting on a malware infection report (see chapter 10), a single scan job for a specific folder can help quickly check whether a malware infection has been completely remediated. In case of a remediated infection, a single scan job of a client's memory and autostart help ensure there are no traces left.

Like periodic scan jobs, single scans can be easily managed from G Data Administrator's TASKS module. The Job scheduling settings are slightly simpler than those for periodic scans. Single scan jobs can be defined in advance, with or without a starting time, and run directly when required. If no time is indicated, the job will not be scheduled for a specific time, but it will be shown in the TASKS overview. By selecting the job and clicking the toolbar button RUN NOW, it can be executed at any moment. Under SETTINGS, the administrator can define several scan parameters. End users can be allowed to halt or cancel the scan job. While this can be useful in case of longer jobs, the nature of the single scan job usually means it should be finished as soon as possible, without any possible user interruption. Similarly, notifying the end user when a virus has been found is often not necessary: when the administrator is monitoring the job, possible malware infections will be quickly acted upon. Optionally, the client can be powered down once the scan job is completed. When further actions are required after the scan, a shutdown is impractical; only if the scan will run as its last task of the day, a client can be shut down automatically. Finally, the client should report its scan progress regularly to the ManagementServer, so the administrator can closely track it.

The SCANNER SETTINGS of a single scan job should be tailored to the circumstances. Since detection

should be optimal, in most cases both scan engines should be used, as well as heuristics and rootkit checks. The types of files and areas that should be scanned depend on the location of the (possible) infection: all files, program files and documents, archives, e-mail archives, system areas: include all locations that can possibly be infected. In some cases, it might be most practical to select an existing full scan job for the affected client(s) and click **RUN NOW** to run it outside of its schedule. The **ANALYSIS SCOPE**, like the settings, should be altered to match the locations of the malware infection. A full hard drive scan is more secure, but may duplicate effort of the full scan. In any case, a memory and autostart scan should be carried out to see if any traces of the malware remain in memory.

9.3. Exceptions

There are many reasons to exclude files or folders from on-demand and idle scans¹⁰. Some files may be wrongly recognized as malware, requiring an exception. A quick scan will, by default, exclude many file types and possibly even low-risk folders. A full scan may exclude archive or database files that will be scanned in a separate scan job. Performance, time, or false positives can all be valid reasons for defining an exception. However, care should be taken when deciding to exclude hard drives, folders, files or file types from a scan job. Almost all file types can potentially contain malware. Every file on each hard drive should be included in at least one periodic scan job, whether it is a quick scan, a full scan or a custom scan job. At the same time, scanning a file more than once between different jobs is not always useful. Alternating periodic scan jobs can pick up on the same file more than once, but having large sections of a client's hard drives scanned multiple times is unnecessary.

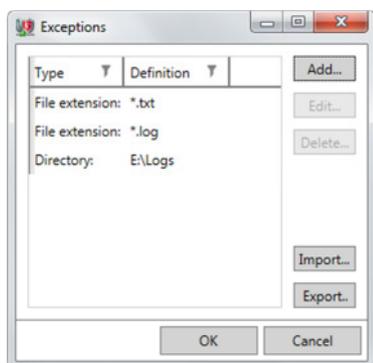


Image 43: G Data Administrator, Client settings module, General, Scan jobs, Exceptions

As with all settings, exceptions apply to the client or group that has been selected in the client management area. It has to be carefully considered for which clients and using which module setting exceptions will be defined, because that delimits the exception scope. Exceptions should be defined for a group that is as small as possible: adding an exception on a client that does not need it opens a window for malware infections. Folders and drives can be defined as an exception that is valid for a specific job only by creating a scan job that do not include them in the analysis scope. Folders that only contain files that do not need to be scanned, can be left out of the regular quick or full scan, but should still be scheduled in a biweekly or monthly job. Alternatively, the exception list on the **GENERAL** tab allows administrators to define file types and folders as global exceptions, which will not be included in any job or in the Idle scan. Because exceptions defined on this list are very wide-reaching, they should only be added in case of a serious problem, such as an essential file being falsely recognized as malware.

¹⁰ For more information about excluding files from real time scans, see chapter 8.2.

10. Handling a malware infection

It does not matter how well end users have been educated about the risks of visiting shady websites or opening e-mail attachments: at some point, malware will reach a client and attempt to infect it. The various layers of the G Data security solution will cooperate in blocking the threat, so no harm will be done to the system or the network. A report will be added to the `REPORTS` module, optionally notifying the administrator(s) automatically¹¹. While malware is blocked completely automatically, it does not mean that a threat can be ignored. Administrators should inform themselves about the threat: where did it come from, what is it capable of doing and what is the risk of encountering it again?

The advantage of having malware blocked automatically is that no immediate action needs to be taken on the client. However, a single malware infection can be the tip of the iceberg: especially corporate networks are often targeted by sophisticated campaigns that consist of multiple attacks using different vectors. G Data Administrator offers excellent logging capabilities and statistics modules to evaluate infections and decide if further measures should be taken. The starting point for any further analysis will be the `REPORTS` module, but many of the G Data security modules can assist in researching a malware infection and prevent further issues.

For organizations that have risk management procedures or similar measures in place, a malware infection can trigger several events. The automated detection and mitigation and extended mitigation stages can be merged with or altered to suit the existing procedures, while the information gathered during the post-infection stage can assist in making sure the attack vector can no longer be abused.

10.1. Automated detection and mitigation

Detecting malware is a fully automated procedure. When one of the security layers of G Data locates a threat, whether during a scheduled scan or as result of a real time security module, it automatically carries out the measure which has been defined in advance and adds a report to the `REPORTS` module. For that reason, it is important to decide the course of action beforehand. Two mitigation paths should be followed: taking care of the infection itself, and preventing further distribution across the network. The former can be configured directly using G Data's `CLIENT SETTINGS` module, while the latter requires a comprehensive overview of network clients and their security settings (see chapter 10.2).

How a malware-infected file should be handled can be defined for each security layer individually. For the real time file system protection, the `MONITOR` tab offers the appropriate options under `REACTION TO INFECTED FILES`. On the same tab, the `BEHAVIOR MONITORING` component can be configured. The e-mail scan can be configured on the `E-MAIL` tab, whereas scheduled scan jobs offer reaction settings on the `SCANNER` tab. Some measures can be combined, such as disinfecting and quarantining a file. For example, this allows administrators to configure the security client to first try to disinfect the file, moving it to the quarantine only if disinfection has not been successful.

10.1.1. Block access

The first action that should be taken once malware is detected is blocking the affected file to prevent it from running and infecting the system. This is the absolute minimum. Not blocking an

¹¹ It is recommended to configure alarms, e-mail messages that are sent to the administrator in case of a malware infection or other trigger event. See chapter 6 for more information about (automated) monitoring.

infected file will allow it to infect the system. Depending on the type of malware, this leads to system instability, performance loss, data loss or worse. All G Data security modules offer the option to block infected files (or halt the process as soon as it exhibits suspicious behavior, in case of the behavior monitoring module). Once a file or process is blocked, the client immediately sends a report to the ManagementServer, where it will trigger the optionally configured notification procedure.

Blocking malware successfully prevents infection, but does not take care of the file itself. The offending file or process will remain on the hard drive and can still try to run at a later stage. This does not form a problem for access by clients that are protected by G Data security modules. However, other clients that have access to the file without being protected, can still get infected. Therefore, it is recommended that there will always be an additional action carried out after blocking the file: either disinfecting, quarantining or deleting it.

10.1.2. Disinfection

Malware can manifest itself as isolated file, solely created to infect systems. In other cases, it consists of one or more components that attach themselves to legitimate files. This can serve several purposes. A user is more likely to open an infected file if it does not look infected. Hiding a malware module in a Word document or third party executable makes it more likely that an end user inadvertently runs in. Furthermore, detection of malware is more difficult when it is wrapped in an innocuous file. Finally, malware that manipulates documents or programs that are already on the client has more leverage over the end user, because it can threaten to remove, hide or encrypt files.

The Disinfection action attempts to remove the malware from the file that it has infected. After a successful cleanup, the file can be used without the risk of infecting the system. For the security modules that offer the option, Disinfection is the recommended setting. When malware is detected, the security module evaluates the file and checks if disinfection is a possible strategy. Not all files can be cleaned: some malware irreparably damages files. Especially difficult is the cleanup of files that have been manipulated by malware, but do not contain malware themselves, such as encrypted documents. For this reason, the real time protection modules focus on preventing infection from happening. Disinfection is a great option to try, but should be combined with a fallback option in case disinfection does not work. For all modules, disinfection can be combined with a different mitigation setting: blocking file access, moving the file to the quarantine, or removing it.

10.1.3. Quarantine

Making use of the Quarantine function is recommended, especially in combination with disinfection. The secured Quarantine folder functions as a safe for infected files. Each security module can move detected malware to the quarantine, where it will be renamed to prevent it from being run. Access to the quarantine is available for administrators and users (if enabled). Every file in the quarantine has its own report in the REPORTS module. The list of reports can be limited to quarantine reports by using the appropriate toolbar button. From the quarantine, the affected file(s) can be cleaned and moved back, or deleted by clicking the toolbar buttons or by right-clicking the report and choosing the appropriate option. When working with quarantined files, it should be made sure that they are not run before they have been disinfecting. If a file cannot be disinfecting, do not attempt to run it anyway, but leave it in the quarantine or remove it instead. The only exception is a false positive: if a file has been wrongly marked as malware, it should be added to the exception list of the affected module, and can be moved back from the quarantine. Extra attention is required when whitelisting suspicious files: accidentally allowing malware back onto a clean client leads to severe system

problems and data loss. Quarantined files can be sent to G Data for additional analysis. In case of a (suspected) virus infection, this helps improve the future detection rate. However, due to the high volume of file submissions, no individual responses can be sent.

10.1.4. Remove file

Removal of an infected file is the most thorough mitigation method. Whenever malware is detected, it is automatically deleted and will not be able to infect the client or to be distributed to other machines. Enterprises that insist on maximum security, no matter the practical consequences, can configure the security modules to immediately delete all infected files. However, this is not the recommended setting. Depending on module and scanner settings, files could incidentally be falsely recognized as malware. Immediately deleting them could then lead to data loss. In all cases, quarantining the file, optionally combined with disinfection, is the recommended setting.

10.2. Extended mitigation

In addition to the automated measures that each security module can carry out, there are several actions an administrator should take. First of all, on the client side, the end user(s) can be informed about the infection. Users may have inadvertently opened an infected website, file or other resource. The virus warning will show them that access to the resource has been denied and that the file has been quarantined or removed.

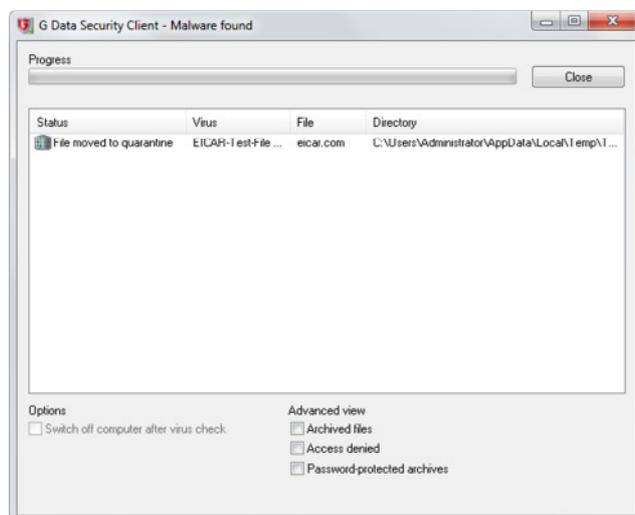


Image 44: G Data Security Client, Malware found

The warning window can be enabled in the CLIENT SETTINGS module by checking NOTIFY USER WHEN A VIRUS HAS BEEN FOUND. It can be useful to notify end users if malware was blocked. They will know to stay away from the affected resource in the future, and might be able to provide the administrator with more details about the infection attempt. However, for inexperienced end users, the prompt may cause confusion or anxiety, leading to support calls. For those users, the virus notification should be disabled. An alternative is manually sending a message to the affected client using the MESSAGES function on the CLIENTS tab. This allows for more customization, sending more appropriate information to the end user. However, this process cannot be automated, so there will be some time between the infection and the moment the administrator manually sends the message.

Similar arguments for and against user intervention apply to the question whether users should be

allowed to open the local quarantine folder. This option can be enabled in the CLIENT SETTINGS module and lets users check a simplified version of the quarantine reports that are normally only available in the REPORTS module. G Data Security Client will offer a QUARANTINE window through its system tray context menu. The Quarantine shows the date and time of infection, the virus name, its filename and the directory where it was found. Users can disinfect, move back or delete files. In case of a false positive, an end user can restore the quarantined file without administrator intervention. At the same time, this shows the danger of the function: files that are possibly still infected could be moved back without any cleaning procedure. It is recommended to keep this option disabled for all clients. If an end user or administrator needs to access the quarantine on a client, the option can be enabled for that client only, making sure it requires a password (enable PROTECT OPTIONS WITH A PASSWORD in the CLIENT SETTINGS module).

On the server side, administrators can carry out several measures once a virus infection is reported. The first thing to do is ensure that a possible infection cannot spread to other network clients. How to do this depends on the severity of the infection. For simple, contained malware infections, taking care of the file by quarantining or deleting it will do. If an outbreak has the potential to infect other network clients, especially those that are not being protected by a security solution, the infected client should be disconnected from the network as soon as possible by (temporarily) revoking its access on network level, or by physically disconnecting it. This does require the administrator to continue infection research and mitigation locally, and is therefore only recommended in acute cases. In any case, scheduling an immediate full scan for the affected client is recommended (see chapter 9.2.2), checking its hard drives, memory and autostart for traces of malware. When using the PatchManager module (see chapter 15.3), the client should be checked for outdated software. Any missing patches should be deployed as quickly as possible. The PolicyManager module helps block applications, devices or web content and can be configured as part of an enterprise content policy. A malware infection can be used as a starting point for a new policy: if the attack vector is known (such as a specific website, an infected USB drive or other removable medium), PolicyManager can block it in future cases (see chapter 14).

10.3. Analysis

After the initial mitigation procedure, the malware infection should be analyzed. This procedure does not involve a strict protocol or pre-defined actions. The focus is obtaining information about the malware to prevent future infections. Due to time constraints, some administrators may not have time for this, while others can research all aspects of the infection and gather facts and advice. Important is one thing: the malware has already been blocked once and will be automatically blocked if it occurs again. Performing additional analysis and gathering more information helps optimize the protection process, but is not required. The more information is found, the more specific the additional measures that the administrator can implement, such as blacklisting specific websites or programs, or configuring additional scans.

Starting point for research of an infection is the Reports module. The file name of the infected file will be listed, in addition to the folder where it was found and the name of the virus. The Reported by column shows which security module has detected the virus, which helps determine how the virus tried to attack the client. For example, reports from the Monitor module refer to files that were written to or read from the file system. Most Monitor detections are triggered by files that are being read from an external device (such as a USB stick), or written by a process (such as a file being downloaded in the browser). Based on the module that reported the virus, extra measures can be planned, such as an extra scheduled scan job, adding a website to the PolicyManager blacklist or

tweaking the Monitor settings.

The virus name is an important entry point in finding information. Both the complete virus name as well as its family or category components can help locate essential tips about dealing with an infection. The Reports module will display the virus name as it was reported by the security module and scan engine. The first part of the name often refers to the type of malware: Trojan, Adware, Generic, or others. This makes clear what the main function of the malware is and may clarify how it infected the system. For some viruses, the malware type will be followed by a platform designation, such as JS (JavaScript) or VBS (VBScript). The platform defines the programming or scripting language that the malware was written in, or is exploiting. While this is valuable information, it cannot directly be used to close attack vectors. Blocking a complete platform from running on client PCs is a drastic measure, which usually affects a large number of legitimate programs. However, looking at the particular security options of a platform and keeping track of its vulnerabilities is recommended practice (even before an infection is detected). In addition to the platform, some virus names contain an explicit reference to the bug that they are exploiting. For example, Exploit.CVE-2011-3402.Gen exploits the vulnerability described in the Common Vulnerabilities and Exposures (CVE) database under entry 2011-3402. CVE database entries can be extremely helpful in finding out information about known vulnerabilities and their mitigation methods. The US Department of Homeland Security maintains a CVE database at <http://nvd.nist.gov> which offers free access to the latest CVE bulletins. The rest of the virus name usually consists of a free form name, adapted from one of the identifying properties of that particular piece of malware. The name often differs between antivirus vendors, but can help kick start an online search for additional information.

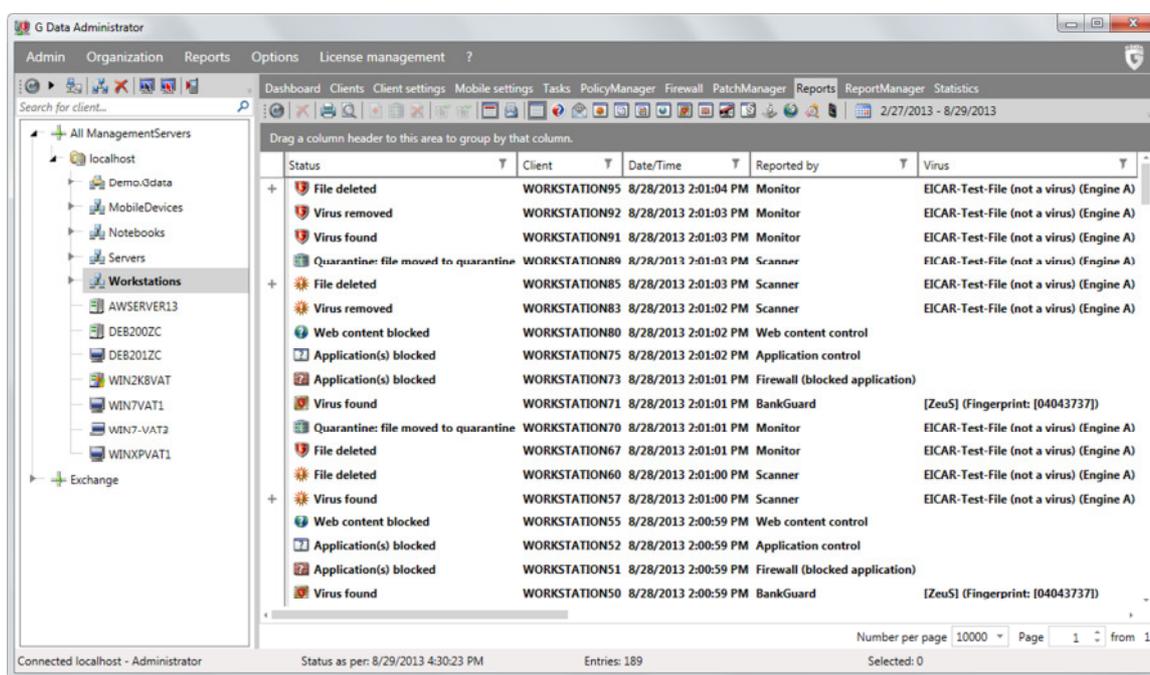


Image 45: G Data Administrator, Reports

11. Mobile device management

Enterprise networks do not consist solely of PC devices. To support increased connectivity for employees, many businesses have started handing out mobile devices, such as tablets and smartphones. Making resources and information available on mobile devices allows employees to be productive outside the traditional work environment, but comes with its own risks. Mobile malware can compromise device security, leading to information theft or damaged equipment. This type of threat is even more applicable to networks allowing unmanaged devices to connect. Enabling access to enterprise resources for employee-owned private devices can be a viable way to increase productivity while saving on device costs, but requires a strict safety policy to prevent malware infections in the corporate network.

Similar to the way in which PC devices are protected, mobile phones and tablets can also make use of an agent-based G Data solution. Mobile device management is included in every G Data solution, and allows administrators to deploy a client app to Android-based devices. The app can be used to centrally schedule web protection and virus scans.

11.1. Managing mobile devices

Mobile devices are managed through the same interface as other clients. G Data Administrator lists mobile devices in its client management area, and configuration takes place using the same tab system other clients use. Mobile devices can be deployed using the appropriate function in the client management area toolbar, and will show up automatically once the client app has been installed and made its first connection to the ManagementServer. This involves setting an installation password in the `SERVER SETTINGS` window; see chapter 4.7.5 for more information on mobile device deployment. Devices are automatically numbered, but their device name can be changed on the `MOBILE SETTINGS` tab. Adding an additional note helps tell the different clients apart. Because mobile devices show up as clients in the list, they can be moved to groups. Creating a dedicated group for mobile devices is recommended, eventually with subgroups for the various departments that may be using managed mobile devices.

Settings for mobile devices can be managed on the various panels of the `MOBILE SETTINGS` tab. As part of device management, update schedule and synchronization should always be defined. Both settings depend on the usage pattern for the device. Devices that are often connected to a wireless network (WLAN) can be configured to update their virus signatures automatically and synchronize data with the ManagementServer every few hours. Devices that are mostly used outside the company network, or connect to the internet using a mobile data plan, can be configured to update less often, or manually, or only when connected via Wi-Fi. The same applies to synchronization: different settings can be configured for Wi-Fi and mobile data plans.

11.2. Real time and on demand protection

Like desktop and laptop clients, mobile clients are also vulnerable to malware infections. Rooted devices in particular do not have sufficient protection mechanisms against malicious apps from unknown sources, but even malevolent apps that manage to sneak their way into the official app stores can have severe implications. Similarly, websites may try to serve malware, take advantage of vulnerabilities in the operating system or otherwise deceive the end user. For this reason, Internet Security for Android offers modules that provide real time as well as on demand protection.

The `GENERAL` tab of the `MOBILE SETTINGS` module contains options for all security modules. Web

protection provides real time protection when using the Android browser. It is recommended to keep this setting enabled. Because web protection can produce a small amount of data traffic, it can be configured to work only when the device is connected via WLAN. The second real time protection component is the automatic virus check, which can be enabled by selecting **AUTOMATICALLY** under **VIRUS CHECK**. Whenever an app is downloaded and launched, the virus check transparently checks it for malware and blocks the installation if it is found to be malicious.

On demand protection is available in the form of a full virus check for the complete device. Under **VIRUS CHECK**, enable **PERIODICALLY** to let Internet Security scan the device regularly. Two types of scans can be configured: **ALL APPLICATIONS** OR **INSTALLED APPLICATIONS**. As long as the device is not slowed down too much by the scan, a regular All application check is recommended to make sure no malware is lingering on storage media (such as an SD card). Depending on how often the device is used and how often new software is installed or saved on it, the interval can be set to 1 day, 3 days, 7 days, 14 days or 30 days. In most cases, it is recommended to perform a daily check: the scan does not cause any noticeable slowdowns, and provides maximum security. To make sure that the virus check is not draining the device battery, it can be configured to only take place while the device is recharging. Alternatively, check **POWER SAVE MODE** to make sure that the scan is postponed if the device is in power save mode. This saves battery, but opens a vulnerability window during which existing malware may go unnoticed.

11.3. Device policies

Technically, deploying mobile devices in a corporate network is no challenge. The workflow for managing a mobile client closely resembles that for a desktop or laptop client. However, mobile devices are often used in corporate as well as private contexts, mixing contacts, communication and applications. Extra attention should be paid to usage scenarios and the policies used to enforce them. For example, administrators could decide to ban rooted devices from the network, allow private devices only if certain requirements are met, or configure all connected devices according to a very strict corporate policy.

Before any policies are defined, mobile devices should be divided into groups. Ideally, all devices in the same group can be managed using the same policy, and all devices using that policy are added to the same group. As with regular clients, there are several types of groups that can be defined, based on network zones and client roles (see chapter 7.1). A logical classification would be to add mobile devices to groups based on corporate departments. If there are usage scenarios transcending department boundaries, grouping devices according to their use is an alternative.

The most effective way to prevent malicious activity from mobile devices is to restrict their access to the corporate network. Without direct access to network resources, mobile devices cannot easily infect other corporate devices with malware, or otherwise execute malicious behavior. On Android devices, the largest threat comes from rooted devices. If the end user has obtained root access to the device, any amount of security on the operating system and app levels can easily be subverted and if malware manages to infect the device, it gains virtually unlimited access to operating system functions. In order to stay in control of managed mobile devices, it is therefore recommended to refuse network access to rooted devices. On the **POLICIES** panel of the **MOBILE SETTINGS** module, the administrator can define the corporate WLAN network's **SSID**, **PASSWORD** and **ENCRYPTION TYPE**. If **ALLOW ROOTED DEVICES** is unchecked, rooted devices that have Internet Security installed will be locked using the remote maintenance password and access to the WLAN will be denied.

In addition to blocking rooted devices, the **POLICIES** tab offers other settings. For each device or

group, a phone type can be defined. For devices that have been issued by the company and are only meant for company use, the phone type Corporate is recommended. This will lock down the client-side setting menus of Internet Security, so that users cannot inadvertently change remotely managed settings while connected to the corporate network. Phone type Private can be used for devices that have not been issued by the company. Phone type Mixed serves company-issued devices that are used for corporate as well as private communication. For each device, the administrator can enable or disable camera access (for devices using Android 4.0 and higher). To protect data stored on the phone, encryption can be mandated. When ENCRYPTION REQUIRED is checked, the device will automatically open Android's encryption settings window to allow the user to enable it. The window cannot be closed until encryption has been enabled.

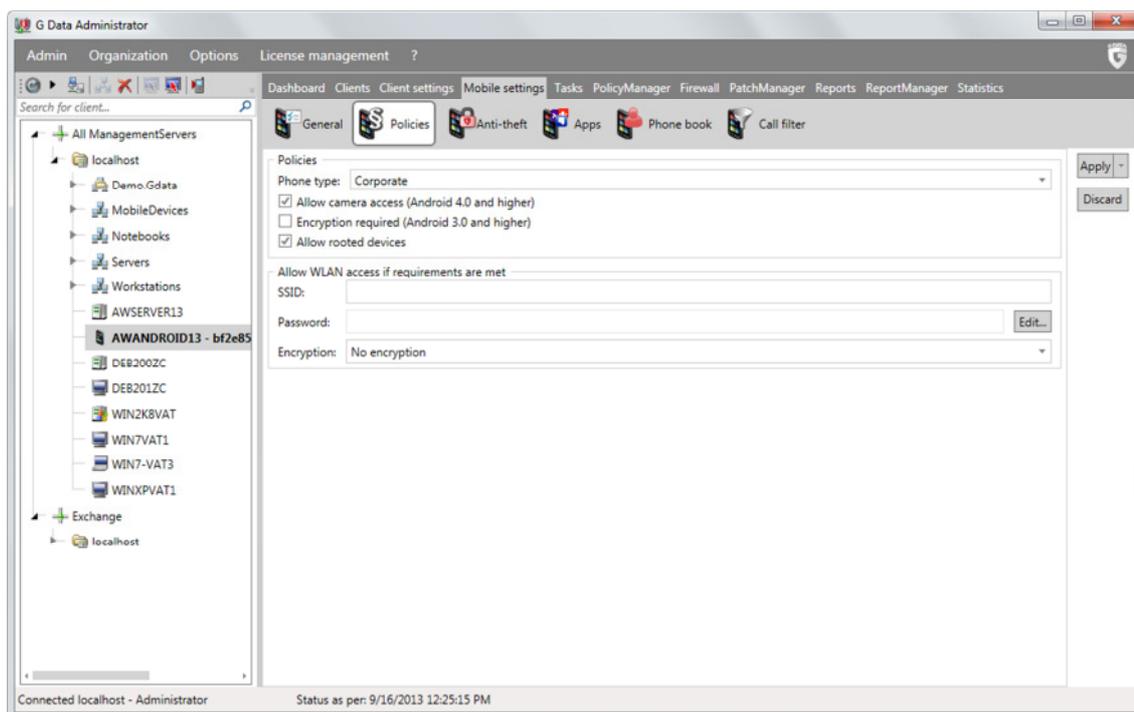


Image 46: G Data Administrator, Mobile settings, Policies

Administrators can enforce other types of controls than just device policies. App usage can be managed using a whitelist or blacklist approach (see chapter 11.4). Calls and SMS messages can be filtered and the phone's contact database can be centrally managed (see chapter 11.5). Whichever kind of policy is defined, end users should be informed of the fact that their device is being managed remotely and that the administrator can decide to take far-reaching actions, such as blocking access to the device (as a corporate policy) or even wiping the device (in the context of anti-theft measures; see chapter 11.6). To this end, the CLIENTS module offers the possibility to have Internet Security display an end user license agreement (EULA) to which the end user must agree. The CLIENTS menu lets administrators write and manage EULAs under EULA MANAGEMENT. Any number of agreements can be created. With a mobile client selected, an EULA can be assigned or removed by choosing EDIT ASSIGNED EULA or REMOVE ASSIGNED EULA from the CLIENTS menu.

11.4. Apps

Part of the attraction of mobile devices is the fact that their default functionality can be expanded by installing apps. Even in a corporate context, this can be extremely useful: productivity tools or configuration apps can significantly increase the number of use cases for mobile devices. At the

same time, corporate devices should provide a controlled environment, making sure that apps cannot cause compatibility problems, obtain sensitive information, or spread malware. App management is a powerful way to control the functionality of a mobile device, balancing security with usability. Administrators should be aware at all times which apps are running on managed devices, and block or allow them as required.

The MOBILE SETTINGS module offers elaborate app management possibilities on the APPS tab. As a first step, it can be used to take an inventory of apps that are in use of mobile devices in the network. Each app is listed with its name, version, size and install date. For each app, administrators should obtain information about its vendor, its functions and its version history, insofar information sources are available. For many apps, the official app store(s) will provide enough details, for others it may be necessary to look up the vendor's homepage. Based on this information and on the intended use of the device (based on the device group and type and the network zone), apps can be added to the whitelist or blacklist. This will allow respectively block the listed apps. Using the defined password, apps are blocked from running.

Whether to use a blacklist or whitelist approach depends to which extent the device should be locked down. If the goal is to only block a few known bad apps, while allowing the user relative freedom, a blacklist approach will do. At the very least, the Android Settings app and Internet Security itself should be password-protected. This will prevent the end user from tampering with any settings. Blacklisting the official app store makes sure that no other apps can be installed. To completely control a device's app experience, the whitelist approach is the most reliable option. Whitelisted apps can be used without any limitations, but all other apps are blocked. This is most useful for devices that are configured for maximum security, or for a single workflow. For example, a device that is only to be used by sales representatives may be run in whitelist mode, allowing only the phone component and the sales database frontend to be used.

11.5. Contact management and filtering

For devices that are used in a corporate context, controlling communication streams can be essential. Blocking apps can help if communication should be entirely prevented, but in some scenarios a more fine-grained filter should be deployed. Rather than completely blocking the Phone app if a device is only meant to be used for work-related communication, outgoing and incoming calls could be filtered if they do not meet corporate criteria. For example, a company that supplies its employees with phones to communicate with headquarters while on the road could block all phone calls except those with pre-approved corporate contacts. To manage contacts on the phone, the Corporate phone book can be used. Even without using any filtering possibilities, blocking the built-in device phone book and populating Internet Security's corporate phone book can be an effective way of ensuring control over contact information.

The PHONE BOOK and CALL FILTER tabs of the MOBILE SETTINGS module together offer extensive contact management and filtering possibilities. The basis of all functionality is the contact database. It functions as a central hub for all corporate contacts, based on which phone books can be created for various devices, as well as targeted call and SMS filters. The database can be opened by opening the PHONE BOOK OR CALL FILTER module and clicking SHOW CONTACT DATABASE. Contacts can be added manually by clicking the appropriate button in the toolbar. The CONTACT window offers input fields for name, address, email/fax/phone and organization details. For organizations with a limited number of contacts, or for small managed phone books, entering contacts manually is a practical way to quickly populate the contact database. If the network is using Active Directory, the IMPORT CONTACTS

button can be used to import contact data. Select the appropriate domain and click OK to import all contacts from the domain. From the CONTACT DATABASE overview list, imported contacts can be edited or removed. Click CLOSE when the database has been populated with all necessary contact information. With all contacts defined in the CONTACT DATABASE, they can be distributed to the appropriate devices. For example, all devices can be supplied with a complete list of colleagues' direct extensions. Alternatively, combined with a block of the standard phone book app and use of the CALL FILTER module, groups of devices can be allowed access only to certain explicitly deployed phone numbers in the Phone book.

The CALL FILTER module can be used for extensive filtering of incoming and outgoing communication. It functions as a filter on the built-in device phone book. Rather than completely blocking the Android phone book app, the filter allows granular control over communication streams. For example, enabling the whitelist mode, no incoming or outgoing calls will be allowed, except for those numbers that have been added to the whitelist. In blacklist mode, communication is generally allowed, but specific numbers can be blocked. Numbers that have been defined in the PHONE BOOK module are always allowed through the filter. Additionally, incoming or outgoing calls from or to numbers in the device phone book can be generally exempt from the filter.

11.6. Anti-theft

Although mobile devices can be a worthwhile addition to enterprise workflows, they do increase risk levels for the physical infrastructure and data-based workflows. Between employees bringing confidential files with them on the road and mobile devices getting lost or stolen, it is easier than ever to accidentally leak confidential files. To make sure that corporate e-mails, documents and other communication cannot be accessed when a device is lost or stolen, several measures can be defined, all achieving different goals. Firstly, it can be helpful to try to recover the device. Locating it using GPS technology or triggering an alarm sound can help. If locating the device is not an option or did not yield any usable results, locking it will make the device useless. As a last measure, devices can be reset to the factory defaults, wiping all data on the device.

The MOBILE SETTINGS module offers access to anti-theft measures. There are several ways in which they can be triggered, automatically as well as manually. To enable all measures, several settings have to be configured. A REMOTE MAINTENANCE PASSWORD (a numerical PIN code) should be entered. It will be used as a password when sending SMS commands, and as a lock screen password if a lock screen password has not been explicitly defined. Enter a TRUSTED PHONE NUMBER to make sure that the password reset command cannot be sent by anyone – only a password reset sent from the trusted phone number will be executed. Finally, enter the EMAIL ADDRESS FOR NOTIFICATIONS, to receive feedback from actions that provide it.

When a device gets lost or stolen, the easiest method of executing an action on it is sending an SMS message to it. Under ALLOWED ACTIONS, you can enable or disable the commands which can be sent to the device. Send an SMS message containing the following command to the device to trigger the respective measure:

Command	Measure
<i>password locate</i>	Locate device. Send an e-mail to the e-mail address entered under EMAIL ADDRESS FOR NOTIFICATIONS.
<i>password wipe</i>	Reset device to factory defaults. All personal data will be wiped.
<i>password ring</i>	Trigger alarm sound. The device will play a ringtone until Internet Security is started.

<i>password</i> mute	Mute device. Mute all ringtones, except the one triggered by the alarm sound option.
<i>password</i> lock	Enable lock screen using the lock screen password. If no lock screen password has been defined, the remote maintenance password will be used.
<i>password</i> set device <i>password: newpassword</i>	Set lock screen password. Sets the password which is used for locking the device. This does not automatically lock the device: be sure to send the lock command afterwards.

The first part of the command (*password*) is the remote maintenance password that has been defined in the MOBILE SETTINGS module. If necessary, it can be reset by SMS. Using the TRUSTED PHONE NUMBER, send the following command: remote password reset: *newpassword*.

If a device is stolen, its SIM card is often removed to prevent the original owner from contacting the device via its phone number. As a countermeasure, you can define actions to be taken automatically when the SIM card is changed. The phone's lock screen can be enabled, making the device inaccessible, and the device can be located. This will send an e-mail to the address defined under Email address for notifications containing GPS coordinates. Be sure to disable these measures if the SIM card needs to be changed, for example when a device is being retired or assigned to another employee.

In addition to SIM- and SMS-based measures, several actions can also be initiated via G Data Administrator. If it becomes necessary to trigger a measure immediately, the EMERGENCY ACTION function can be used. The device does not need to be connected to the ManagementServer network for this to work: it relies on Google Cloud Messaging, an online service from Google which lets you send commands to Android devices. Google Cloud Messaging does require some configuration. You will need to register a Google APIs account to obtain an API key and sender ID. A Google APIs account can be created at <https://code.google.com/apis/console/> (you need to be logged in with a new or existing Google account). Click CREATE PROJECT to create a new API project. On the SERVICES tab, enable the option GOOGLE CLOUD MESSAGING FOR ANDROID, which requires you to agree with Google's Terms of service. After enabling the option, click API ACCESS in the menu bar to open the API access page. Click CREATE NEW SERVER KEY to create an API key for use on your server. In G Data Administrator, this key should be entered under SERVER SETTINGS > MOBILE. The Sender ID can be found in the browser's address bar. While logged in to Google APIs, you will be browsing URLs with a structure like <https://code.google.com/apis/console/#project:123456789012:overview>. The number (in this example 123456789012) is the sender ID – enter it in the MOBILE window as well.

With Sender ID and API key configured, it is easy to trigger an emergency action from the G Data Administrator interface. Select the device on which you would like to trigger the measure in the client management area, and select the appropriate action. By clicking EXECUTE ACTION, the command is sent to the device and carried out immediately. The commands function identically to the ones that can be triggered by SMS.

Because the anti-theft actions can severely affect the usability of the phone (e.g. by removing data from it), it is recommended that the end user is informed by deploying an EULA. See chapter 11.3 for more information about EULAs.

12. Backups

For many enterprises, the digitization of work processes has appointed computers to be organization-wide information carriers. While access to digitized data is far easier than having to regularly access paper archives, it brings its own set of challenges. Data security and data integrity have to be monitored carefully. Physical problems such as a power failure or a broken hard disk can severely affect workflows, especially if files cannot be recovered. Malware may encrypt, infect or remove files, or documents are removed through human error. To minimize recovery time after a data incident, it is important to plan file backups and recovery carefully. Regularly making a backup by copying files to a secure location makes sure that files are never irrevocably lost. As with infection mitigation, planning ahead is key. If files are lost due to a physical problem or human error, they should be recovered without delay. Making sure that files are backed up regularly by formalizing a backup and recovery plan is essential in guaranteeing business continuity.

A backup and recovery plan is part prevention, part recovery. As a starting point, it can be helpful to classify all data sources on the network according to properties such as value, risk or backup effort. If the company employs information workers, who spend most of their time working on digital documents, those documents should always be secured. Similarly, central data storages such as databases, e-mail servers or collaboration environments should be backed up to make sure that no data are lost in case of emergency. In contrast, some data are interchangeable. A client's operating system usually does not need to be backed up, because it can be reinstalled if its data get compromised. The same goes for software packages for which the installation media are still available. Operating system configuration and software settings, however, may be among the data that should be backed up, if extensive customization has been carried out after deployment. For every decision about the data to be included in backups, storage has to be available. Especially when data sets are regularly backed up, or if they include large files, storage requirements can increase exponentially.

When the data to be backed up have been selected, the question is how often the backup should be carried out. Backup jobs can be planned to be carried out only once, or according to a schedule. A one-time backup is only useful in very specific circumstances, for example when a data set needs to be secured immediately, outside of the defined schedule. In most cases, backup should be scheduled. A recent copy of the data to be secured should be available at all times, to be able to quickly recover in case of emergency. A scheduled backup takes care of this. After configuring it once, there will always be a recent data copy. All that needs to be done is regularly checking logs to see if the backup was carried out successfully. There is a trade-off between security and performance. The more often backups are made, the less data will be lost if a client hard disk crashes or if files are infected by malware. However, making a backup requires time and performance (disk activity) on the client, and disk space on the backup target (server). This issue is mostly addressed by using differential backups (see chapter 12.2), but it is important to verify that the backup target has enough disk space available at all times.

Since backups are an essential part of a security policy, it should be made sure that they are carried out successfully. As with most IT related tasks, it is recommended to appoint a person or team to the task of planning, managing and carrying out backups. Having a responsible person or team will allow for swifter decision making in case of emergency. If an administrator or end user reports data loss, the workflow should be optimized for speedy choices and effective data recovery.

12.1. Managing backups

The G Data Backup module is available as part of the Enterprise solutions and SmallBusiness Security. Its seamless integration allows administration through G Data Administrator's Tasks and Reports modules, and makes sure that backup configuration and management can be carried out with ease. Backup jobs can be planned and managed on the TASKS module tab. It lists all planned jobs – click the appropriate toolbar button to limit the list to backup jobs only. As with every module, all jobs apply to the client or group selected in the client management area. The list columns show the most important properties of the jobs and can be sorted by client, group, status, last execution, interval, scope, or name. For client jobs, the client for which they have been defined is displayed (for groups, the group name). The STATUS column shows the current status of the job. For group jobs, the status can be checked per client by selecting the appropriate client on the left. The INTERVAL column displays the defined backup interval, such as Once for a single scan job, or Daily for a periodic scan job that runs every day. Finally, the SCOPE displays the backup scope that has been defined for the job. When a backup job has been previously run, expanding it reveals a list of status messages for each run. Double click on the status to open a detailed log.

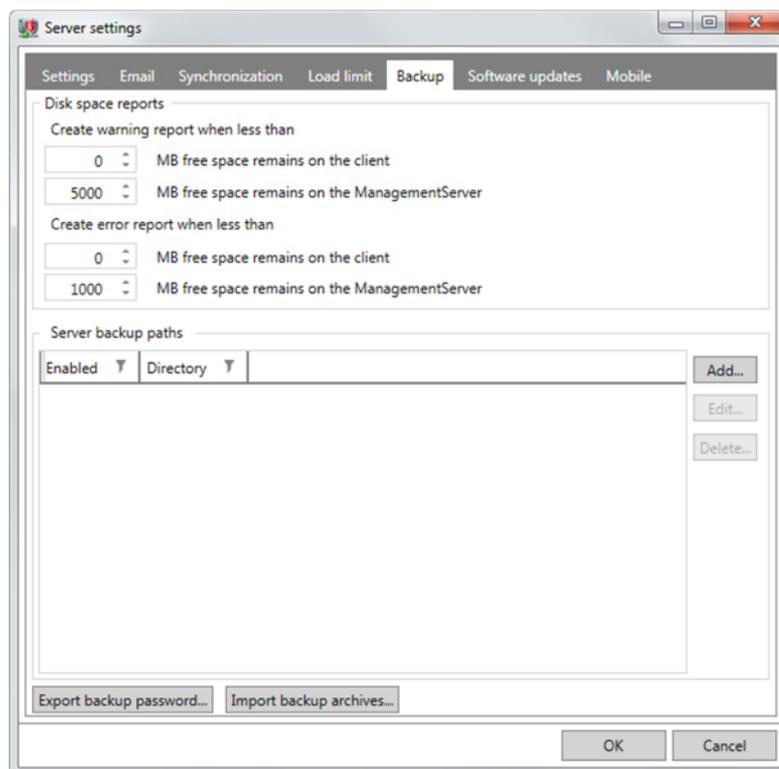


Image 47: G Data Administrator, Server settings, Backup

Backup and restore jobs have individual settings, but there are also a few general options. Most important is setting the storage location. If a large number of backups is scheduled, or more importantly, if they are run often, a large amount of disk space is required. Setting a storage location therefore requires careful planning: storage should be ample and needs to be easily expandable. There are other requirements to storage space. For performance reasons, the target should not be compressed or encrypted. The backup software already automatically encrypts and compresses data – compressing or encrypting them again on the operating system or hardware level decreases performance while not adding extra security. The backup target should be a disk dedicated to storing backups, not containing any other files. This allows administrators to easily archive disks

physically if required by a company's data retention policy, and prevents confusion between regular files and backups. The default backup target is the server folder `C:\ProgramData\G DATA\AntiVirus ManagementServer\Backup` (from Windows Vista/Windows Server 2008 onwards) or `C:\Documents and Settings\All Users\Application Data\G DATA\AntiVirus ManagementServer\Backup` (Windows XP/Windows Server 2003). Because these folders are located on the server's system disk, it is not recommended to use them as a backup target. At least one backup path should be configured, for example another disk in the same machine. When defining multiple paths, G Data Administrator will automatically save the backup file in the first path that has enough free disk space for the backup file. Additionally, this makes sure the backup target can be easily switched by disabling a path (for example, if there is no disk space left).

In addition to the backup targets, the `SERVER SETTINGS` window offers the possibility to export the backup password. This is recommended, since backup archives are password-protected. When importing backup archives, the password has to be entered or imported from a file, so it should be made sure that the password is exported and saved in a secure location. To import an existing backup archive, click `IMPORT BACKUP ARCHIVES`. After a folder containing backups has been selected, G Data Administrator will import its contents and list the backup file(s) in the `TASKS` lists under the header `IMPORTED ARCHIVES`. The backups can be restored like any others, using the `RESTORE JOB` function (see chapter 12.3).

When carrying out a backup job, the client saves the backup in a local cache while it is being transferred to the ManagementServer. When the option `USE CLIENT STANDARD PATH` is checked, the backup is cached on the partition containing the most free disk space. If that is the system disk, the folder will be `\ProgramData\G Data\Backup`; if that is another disk, the directory will be `\G Data\Backup`. By unchecking the option, a non-standard directory can be configured. This is especially useful to force the cache to be saved on a non-system disk.

To carry out a backup, both client and server need to have enough free disk space available, for the backup cache and backup storage respectively. The `SERVER SETTINGS` window lets you configure threshold values for disk space. When the amount of free disk space on client or server drops below the warning threshold, a warning message will be added to the `REPORTS` module and the client cache will be cleaned up, removing every cached backup that has already been transmitted to the server, except for the latest one. When the amount of free disk space on the client or the server drops below the error threshold, an error message will be added to the `REPORTS` module. Server backup storage and client cache will be automatically cleaned up. If there is still not enough free disk space on the server, backups will not be carried out. The error threshold for client and server should cover the size of one or more of the largest backup jobs that will be carried out, with a margin. The warning threshold should be set higher, so that the administrator is timely notified of impending storage problems.

12.2. Create a backup

As with scan jobs, it is recommended to think of the backup schedule as a whole, instead of planning several unrelated backups. Together, all backup jobs should ensure that documents, settings and other important files are regularly copied to a safe location. But an effective backup schedule depends on more than that. Data consolidation, client settings and server performance strongly influence the manageability of backup jobs. Before planning any backups, it is recommended to consider these issues and create a unified backup plan.

Depending on the number of files to be backed up, a backup job can take several hours to complete.

When planning a full backup, make sure to schedule it at a time where the client is not in used, to avoid performance loss. Subsequent partial backups will not involve backing up as much data as the first time, because only new and updated files will be backed up. In that case, scheduling the backup during regular work hours is not problematic. Especially for recurring jobs, making sure the backup moment does not interfere with anything is vital. Other backup jobs, patch management jobs or scan jobs should not be run concurrently with a planned backup to avoid bringing the system to a halt. When planning a backup for laptop clients, the backup can be postponed when the laptop is running on battery power to avoid burdening the hard drive and draining battery power. It will resume as soon as the laptop is plugged in. When planning a backup of a small number of files, which can be carried out relatively quickly, data security can trump user experience and laptop performance, but in most cases, postponing the backup until the next charge is safe.

Selecting the files to backup is the hardest part of defining a backup job. To ease the process, it is helpful to make sure data sources are consolidated, both on a client level as well as network-wide. Rather than having to check each client to see if and where documents and settings are stored, saving them in a standardized local folder or on a server share will save time and makes sure that nothing is accidentally left out of the backup. By default, Windows uses standardized local folders for documents and settings. The `BACKUP` module offers the option to automatically include all of these user folders in a backup job. Even though this is a form of data consolidation, it is not recommended to make a backup of just the user folders: it may include unnecessary files and waste backup disk space, or miss out on documents that were manually saved elsewhere or in an inaccessible user folder. End user file storage is most easily managed if it is located on a central (file) server. Offering a user-based profile folder on a network share as a default save target for documents will allow administrators to secure files for all end users at once by planning a backup for just the network folder. For smaller networks that may not have a file server, Windows' local consolidation can be useful, but administrators should be aware of its limitations, and make sure that all vital files are indeed successfully backed up.

Aside from the server-side possibility of excluding temporary files from backups, local measures can be used as well to minimize clutter and save disk space. When adding a folder to a backup job, administrators should verify that it only contains the type of files that should be backed up – temporary files, generic system files and other low-priority files should be excluded from the backup, moved to another folder or removed.

The temptation to plan one backup job that includes everything should be resisted. All-inclusive backup jobs waste disk space, take a long time and complicate swift, targeted recovery. Backups can be organized by file type, risk or importance, or a combination thereof. For example, important database files should be backed up regularly, while log files on the same machine can be secured with a larger interval. It can be helpful to take inventory of the types of files that are being produced and used on network clients, how important they are, and how high the risk of data loss is. For each group of files, a separate backup should be planned, its execution interval decreasing for important or high-risk files. See chapter 12.2.1 and further for suggestions on backup categories, but be sure to include any file that is important to the organization and should not be lost.

During the backup process, all files that are copied are checked for malware. While the archive is being built, it is made sure that no infected files are backed up. This removes the need to plan an additional scan job right before the backup is carried out. To prevent unauthorized access to the files, all backup archives are encrypted, so they cannot simply be opened from the hard drive. This also prevents malware from infecting backup archives. When the backup process finishes archiving all files, the archive integrity is verified to ensure no problems have occurred while building it. The

file is then moved from the local cache to the server, which will move it to its final backup target folder.

After a backup completes, its log will be available through the appropriate job entry in the TASKS module. For each iteration, whether it is a full or a partial backup, a separate log will be created. Files that are in use at the time of backup may not be copied correctly, or disk read errors can cause complications. Therefore it is recommended to regularly check the logs for each backup to see if it was carried out successfully, and if any additional backups or measures need to be planned.

Additionally, a regular test restore should be planned. To make sure that the backup can be restored correctly, and contains the correct files, a test restore should be carried out to restore the backup to a client or folder of choice. Because the presence of a backup is vital if problems occur on a client, regular test runs or other ways to verify integrity of the backup files are important.

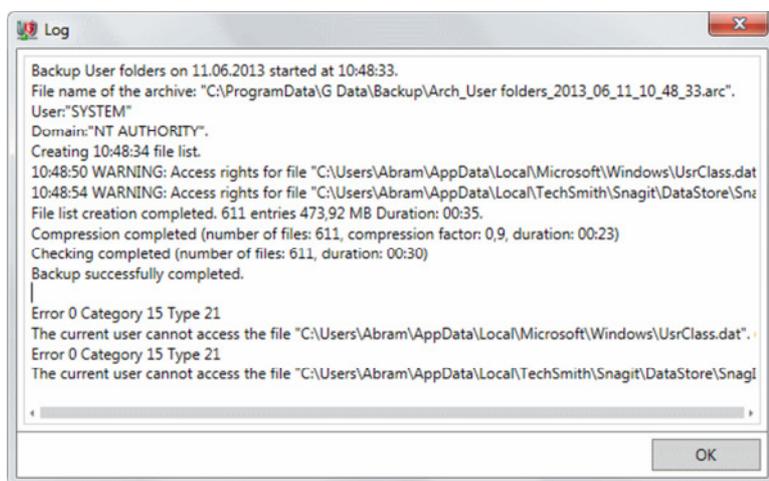


Image 48: G Data Administrator, Tasks, Backup log

12.2.1. Documents

Arguably the most important type of files to back up is documents. In most organizations, information workers spend a considerable amount of time on producing digital files, such as documents, programs, graphics and other assets. These files often form the core of the production process, and any loss of data could set production back or even directly cause revenue loss. Making sure that those files are safe should be the primary aim of the backup plan.

Documents can be located on clients' hard drives or other local or external storage. This can greatly complicate backups: if there is no centralized document storage, a separate backup job will have to be defined for each client, explicitly covering its local document storage folders. For clients that make use of user profile folders, such as Desktop or My Document, a backup job can be defined that covers all local user directories. This will pick up all documents saved in those folders, as well as select settings. Alternatively, a backup job can be created that includes the complete hard drive, but excludes non-document file types (to be defined by extension, such as .exe, .dll or .com). However, the most reliable backup strategy is to control file storage at a much earlier stage. Central file storage, on a network file server for example, allows for a very simple backup job that simply covers that folder tree. By preventing fragmentation at this stage, the chance of missing out on files during the backup stage is greatly reduced.

Job scheduling for document backups depends on the frequency with which the clients are used. In most cases, the client will be used daily, so a daily backup should be performed. Only sporadically,

the backup frequency should be reduced to once every two or three days. The first backup should be a full backup, carried out at a specific point in time when the client is not in use. The full backup will consist of all files, and therefore be relatively large. Subsequent backups can be configured as partial backups. This type of backup (also known as differential backup) saves only the files that have been changed since the last full backup. This saves a significant amount of time in making the backup, and uses less disk space on the backup targets. Restoring the backup will take slightly longer, because it has to be rebuilt from multiple files (the original full backup, and the latest partial backup). If this is a problem, the subsequent daily backups can also be full backups. However, the required storage space will quickly escalate. In that case, configuring a weekly full backup and daily partial backup is the recommended scenario.

12.2.2. Databases

In many IT-powered organizations, certain types of information are stored centrally. This can range from a centrally managed workflow environment to contact databases, sales information or other assets concerning the organization's production. Another important database is the mail server's message storage. In case of a disk failure, the single database containing all of organization's critical e-mails, contacts or product information could be wiped out. To ensure data security, backups should be configured to make sure that the complete database is regularly backed up to a safe location.

Backing up databases can be complicated. Depending on the type of database, its files may be in use continuously. An e-mail database, for example, will not be accessible to the backup process for as long as the e-mail server is running. At the time of backup, make sure that the database files are not in use. Shut down any running processes that make use of the database, as well as background services. For many types of databases this may cause a significant service disruption. In those cases, a backup should be carried out during an existing maintenance window, for example on the weekend or during the night, when there is little to no activity.

How often a database backup should be carried out, depends on how the databases are being used. For vital information that is updated regularly or constantly, such as e-mail databases, a daily backup is recommended. The first backup is a full backup; subsequent ones should be partial backups, to prevent backup storage space from filling up quickly. Databases that are less regularly used can be backed up weekly. In this case, it is easier to plan the backup: there will be less service interruption if a weekly maintenance window can be used.

Some databases can be easily backed up with a file-based backup solution like G Data's backup module, while others have built-in backup tools. Whether you are using G Data's backup module or a built-in database backup tool, it is vital that backups are made regularly and that their integrity is verified to make sure they can be successfully restored if the need arises.

12.2.3. Configuration

Over time, a client's configuration can start to significantly differ from how it was initially rolled out. End users often change their preferences in the operating system and third party software, or administrators deploy changed configurations to take care of local issues. Losing configuration settings can cause annoyance when a significant amount of time is required to reconfigure a system after a disk failure. Making backups of configuration files helps reduce the time needed to reconfigure a system after a reinstall.

Locating all configuration files on a client can be tricky. Some settings are saved in files (typically ending in .ini), others are saved in the Windows Registry, which cannot easily be backed up. For software that supports it, saving settings centrally (either in a central configuration file or using a method such as Microsoft group policies) can prevent a lengthy process to identify the correct files and simplifies backup. In addition, this makes it easier to roll out configuration changes and default settings for multiple clients at once. For smaller networks, administrators may choose to manually identify which software is currently installed on clients, and where its configuration is saved. The Software inventory (CLIENTS module) can assist here in finding information about software on the clients, but it will require a considerable amount of time to locate the configuration files. For those that can be located, a backup can be planned. A backup frequency of once a month will cover most situations, but for clients that are reconfigured more regularly, a weekly or daily backup is also a possibility.

12.2.4. System files

Backups do not only need to focus on files that were generated by an end user. Executable files and other components that belong to an operating system or third party program can be backed up as well. As with database backups, the process can be difficult, as many system files are continuously in use. An additional problem is the fact that a software backup often cannot function on its own. Backing up its Program Files folder does not take into account the many dependencies, such as other software, frameworks or certain operating system features or updates. Restoring such a backup to an empty system does not guarantee that the software will work.

To make backups of software, a better solution is to safeguard the original installers and save them on a server. As for the operating system itself, the restore process can be simplified by making sure there is always a full system image available of the operating system with its latest updates and all the third party software that is typically used. This image resembles a backup, but does not include any user documents. Instead, it is built from a 'clean' operating installation with its latest updates, and all required software. Keeping this image up to date requires some effort, but greatly reduces the amount of time needed to restore a full system. This works best in a scenario where clients do not save any personal data locally, but on a file server instead: a client disk failure does not cause any data loss, and an operating system image can quickly be restored once a new hard disk is in place.

12.2.5. Backup on demand

Backups can be carried out manually, instead of according to a schedule. This can be useful for scenarios where a client is typically not backed up at all, due to its low importance or low risk of data loss. Alternatively, if an impending hardware failure seems to threaten a client's stability, a backup on demand can quickly save data before the client fails. This type of backup jobs can be planned ahead, to enable a quick response time in case of emergency. However, as a job cannot be assigned to another client than it was initially planned for, there is no way to prepare a single emergency backup job. The BACKUP JOB window is fairly simple and allows administrators to quickly define the necessary measures if a backup on demand should be planned.

12.3. Restore a backup

When disaster strikes, it is important to have a recent backup available. When the backup jobs have been properly scheduled, data loss will be minimal. Restoring data is relatively simple, compared to

planning the backups. The minimum requirement is a working installation of the operating system and G Data Security Client (which both may have to be reinstalled, in case of a complete disk failure).

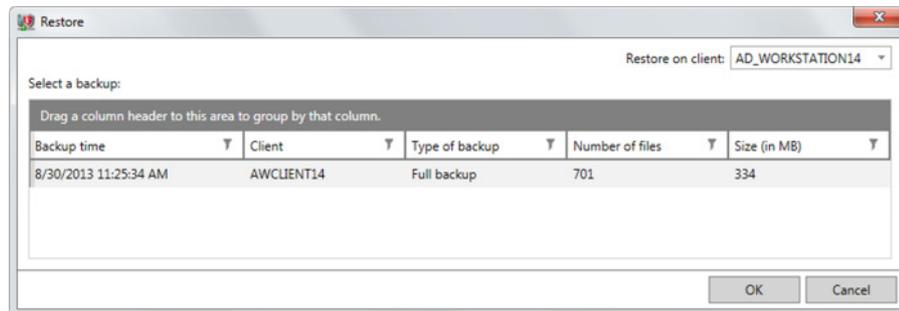


Image 49: G Data Administrator, Tasks, Restore backup

Select the client in the client management area and open the TASKS module to view all existing backup jobs for the client. Right click on a job and click RESTORE BACKUP to open the RESTORE window. Alternatively, select the ManagementServer, open the TASKS menu and choose NEW > RESTORE JOB to get an overview of available backups across all clients. Select a backup and click OK to open the RESTORE SETTINGS window. Upon confirming the recovery settings, a Restore job will be added to the TASKS module. It will be carried out immediately.

If storage is limited, administrators can manually move backup archives to different locations. To be able to restore them, however, they need to be reimported into G Data Administrator. Under OPTIONS, SERVER SETTINGS, BACKUP, the IMPORT BACKUP ARCHIVES button will let administrators import backup archive files from any folder. The backup password will have to be entered, which is available through the EXPORT BACKUP PASSWORD button on the ManagementServer that created the backup. The imported backups will show up in the backups list in the TASKS module.

Restore jobs can be planned to restore a system to the state in which it was when the backup was made. This applies mainly to backups of configuration or system files. In this case, all files should be restored to their original folders, overwriting existing files if necessary. Databases and documents, on the other hand, can be selectively restored, to provide the client only with the necessary files. The FILE SELECTION tab can be used to browse through the backup archive and include or exclude any files or folders. In case a file has been accidentally deleted, a backup archive can be used to locate the most recent backup of that file alone and restore it.

A backup does not necessarily need to be restored to the client it was created from. Any backup can be restored to any client, provided it has enough free disk space. With the appropriate backup selected, the RESTORE window offers the RESTORE ON CLIENT dropdown through which any client can be selected. In the following window, the files that should be restored to that particular client can be selected. This allows administrators to create a backup of a client with one network role (including, for example, more types of files or data sets than other clients), and still be able to use to restore to it to clients in other roles.

It is recommended to test backup archives after the first backup run, to see if the backup settings led to all files being included correctly, and if the backup archive functions. Right click on the backup job and choose RESTORE BACKUP to open the RESTORE window. Verify that the backup archive is listed and click OK to open the RESTORE SETTINGS window. To verify if all files were there, the file and folder list can be used. To check if the backup can be restored without problems, configure it to restore all files and define a new target directory on the OPTIONS tab.

13. Firewall

The firewall is an essential component of network security. As the first line of defense, firewalls filter inbound and outbound data traffic. This ensures that attackers cannot remotely access services that are running on a machine, and that software on the machine cannot contact outside servers. Through the use of a whitelist (firewall rules), certain outside visitors can be granted access, and certain local software can be allowed to contact outside servers. Depending on the network layout (see chapter 1), a firewall can be a hardware device or a software solution. A physical firewall filters all traffic coming in and going out of the network. Many routers, enterprise-grade as well as low-end ones, have a built-in firewall. Software solutions, on the other hand, can be installed on a server or client and provide protection on the machine level. Each network should at least have one firewall active that keeps unwanted traffic out. For all enterprise networks, enabling a firewall as the first layer of network security is recommended. Additionally, clients should be provided with a software firewall. This allows for a granular control over application permissions. Instead of blocking certain traffic for the whole network, in- and outbound network traffic can be managed by defining rules for individual clients or network zones.

Firewall rules should be coordinated centrally, and be configured according to the company's security policies. Across all deployed firewalls, whether hardware- or software-based, rules should be harmonized to prevent conflicts. The firewall rules serve two purposes: managing network security and maintaining company policies (that may not have anything to do with security). For inbound traffic, defining firewall rules is relatively easy: all traffic should be dropped unless it was requested by an end user or program (for example, surfing the web or initiating a peer to peer download). It is purely a security issue, where non-requested traffic is considered to be a malicious connection attempt. For outbound traffic, the rules get more complicated. Some software and services can be denied outbound traffic because they may form a security risk. Others are not malicious, but do not comply with company policy. For example, outbound traffic for a chat program is not malicious, but can still be blocked if company policy states that chats are not allowed.

The centrally managed client firewall rules are enforced by the firewall module of G Data Security Client. It filters all network traffic that is coming in and going out, making sure that unauthorized communication is blocked. G Data Firewall uses stateful inspection to filter network traffic. This method does not only compare single packets to a predefined set of rules, but also considers the previous packets that were exchanged with the same server or computer. This way, packets that are transmitted in an existing TCP stream or UDP connection do not need to be completely examined again, but can be filtered based on the state of the connection (such as IP addresses or ports that have been used previously). This reduces the amount of processing power needed per packet, and makes sure that administrators and end users can quickly define rules to allow a type of connection, without having to worry about the exact IP address, port number, or packet direction.

The firewall module included in G Data Security Client is meant to be used on client machines. While the client and all its protection modules (such as the file system monitor and the firewall) can be installed on servers without problems, the predefined rule sets are aimed at client machines. For use on a server, administrators have to make sure that an appropriate rule set is defined.

13.1. Managing firewall clients

G Data Firewall can be managed through its own module in G Data Administrator. The FIREWALL tab shows all relevant options, applicable to the clients that have been selected in the client management area. When deploying the firewall for the first time, or configuring client settings, it

should be made sure that the firewall has indeed been installed on the appropriate clients (either at deployment time or afterwards). The column G DATA FIREWALL on the OVERVIEW panel shows the firewall status for all selected clients. For clients that do not yet have the firewall component installed, the column will show NOT INSTALLED. For an installation afterwards, the easiest way is using G Data Administrator's remote installation capability. Open the FIREWALL module and select the client(s) on which G Data Firewall should be installed. Right-click the client(s) and choose INSTALL G DATA FIREWALL. This will initiate a remote installation of the firewall component. Like the remote installation of G Data Security Client, this assumes that several system requirements have been met (see chapter 4.7.1). If remote installation is not an option, a client installation package or local installation can be used to install the firewall (see chapters 4.7.2 and 4.7.3).

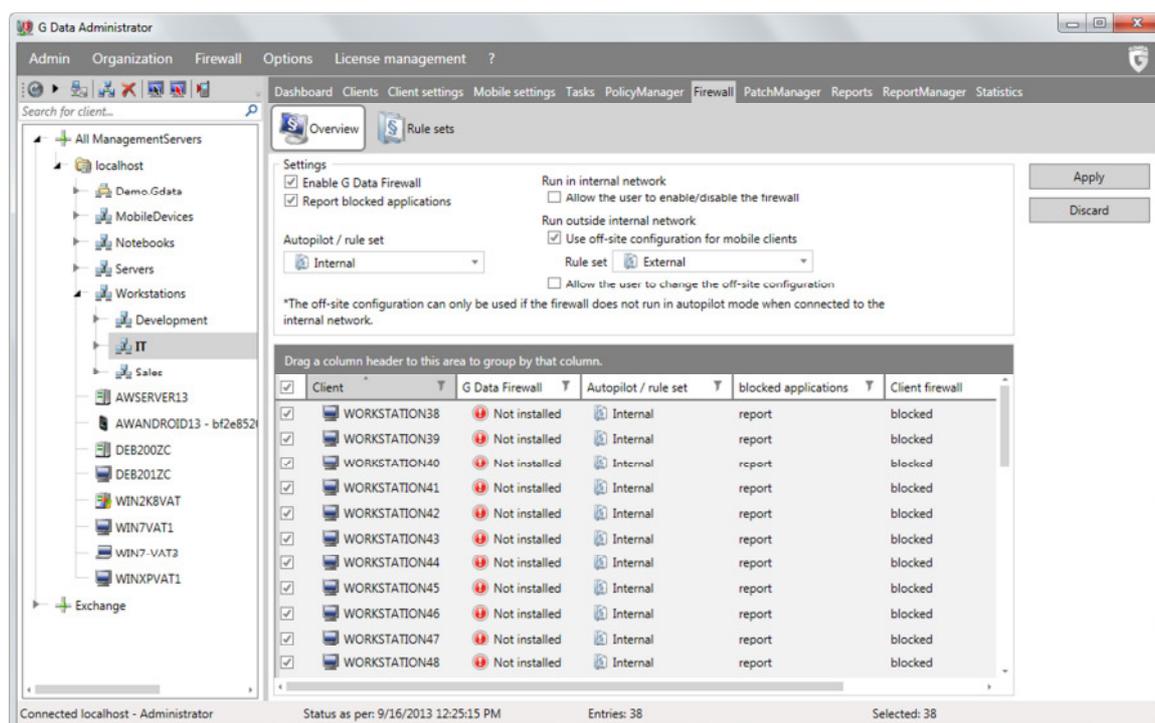


Image 50: G Data Administrator, Firewall, Overview

As with other modules, firewall settings can be applied to one or more clients at the same time. Select the relevant clients in the client management area to have them show up on the OVERVIEW panel of the FIREWALL module. Per client, the firewall can be enabled or disabled. It is generally recommended to enable G Data Firewall at least on all clients that connect to the internet, and ideally on all clients in the network. Only if an alternative client firewall is being used should G Data Firewall not be enabled, in order to prevent conflicts. To keep an eye on client activity, enable REPORT BLOCKED APPLICATIONS. This will make sure that every time an application is blocked, G Data Firewall sends a report to the ManagementServer, which will be displayed in the REPORTS module. If a client application gets blocked by one of the firewall rules, the report will show the offending application and allow administrators to directly add it to one of the rule sets if necessary. As with virus alerts, repeated firewall reports may indicate that an end user is misusing a machine, or that a part of his job cannot be carried out because of usability limitations.

Per client or group, an off-site rule set can be enabled. Enabling a specific off-site configuration is particularly useful for clients that are regularly used outside the enterprise network. G Data Firewall will make use of the normal configuration as long as the client is connected to the enterprise network, but will switch to the off-site configuration as soon as the machine connects to another

network. This ensures security policy compliance for the enterprise network, while allowing an end user more flexibility when the machine is used elsewhere. An off-site rule set can only be enabled if the client is using a rule set for the enterprise network; using the autopilot mode prohibits the use of off-site configuration.

13.2. Autopilot

Autopilot is G Data Firewall's default setting for all clients. This option configures the firewall to carry out its tasks completely in the background. End users will not be confronted with any prompts and administrators will only have to carry out a minimum number of management tasks. Autopilot is very light on management, and can be deployed in networks where only a minimal number of third party applications is used, or where the software inventory is relatively stable. Inbound and outbound connections are automatically evaluated and allowed or blocked. If software tries to establish an outbound connection, the firewall allows the connection if the process is not detected as malware. For the first two outbound connection attempts, a temporary rule allowing the traffic through the firewall is established. If the same process attempts to open an outbound connection for the third time, a rule is added to the autopilot rule set to permanently allow the software. Inbound connections are always dropped, unless they are part of communication initiated by an outbound connection attempt of one of the processes on the system.

By keeping an eye on the `REPORTS` module, administrators can check how often the firewall blocks applications. If it turns out that the autopilot mode prohibits essential outbound or inbound connection attempts, switching affected clients to the manual rule set mode is recommended. The same goes for laptop clients: machines that are often used outside the enterprise network, should be switched to manual rule sets to allow for off-site configuration capabilities.

13.3. Rule sets

When not using the autopilot mode, administrators have a wide range of possibilities to configure the firewall. G Data Firewall makes use of rule sets: collections of application-, protocol- and port-based rules that control the network data flow from and to clients. Rule sets allow for fine-grained control over network traffic, but require more time to be configured than autopilot mode. Administrators should be familiar with network layout, protocols and applications before attempting to configure the firewall using a custom rule set. After a rule set has been created, it can be assigned to one or more clients by opening the `OVERVIEW` panel and configuring it as regular or off-site rule set.

The `RULE SETS` panel is used to create and edit rule sets for the entire network. At the top of the panel, the `RULE SET` dropdown list shows all rule sets that have been defined. By default, the firewall is operating in autopilot mode without any defined rule sets. Click `NEW` to add a new rule set. Creating a new rule set requires entering a name. Optionally, a note can be added to further describe the set (for example, for which clients it is meant, or if any specific applications, protocols or parts are included). Check `STEALTH MODE ENABLED` to prohibit clients to respond to port probes, further increasing security. In the following popup window, firewall rules can be selected from the default rule set. This set contains rules for many common applications, including Windows- and Microsoft-specific functions, but also third party software such as Adobe Reader and Mozilla Firefox. Selecting rules for applications that are in use in the network saves time, but it should be ensured that only the most necessary communication is allowed. Clicking `OK` creates the new rule set, including the chosen rules. Next to the rule set list, the `EDIT` button allows for editing of name and note. The `IMPORT` and

EXPORT button can be used to save and import rule sets, useful for rule set backups or for predefined rules to be added easily.

Rule sets contain an arbitrary number of rules, either pre-defined, custom, or created in response to a report. Connections are evaluated against every rule in the rule set. Rules are listed in order of priority (rank): for each connection attempt, the rule at rank 1 trumps the rule at rank 2, which trumps the rule at rank 3, etcetera. This allows for granular control of ports and protocols, for example to allow UDP traffic on port 2000 in rule 1 and dropping all other UDP traffic in rule 2. Using the RANK control to the right of the list, rules can be moved to other positions on the list. To disable a rule within a rule set, uncheck the checkbox next to its name or edit the rule and uncheck the option RULE ENABLED.

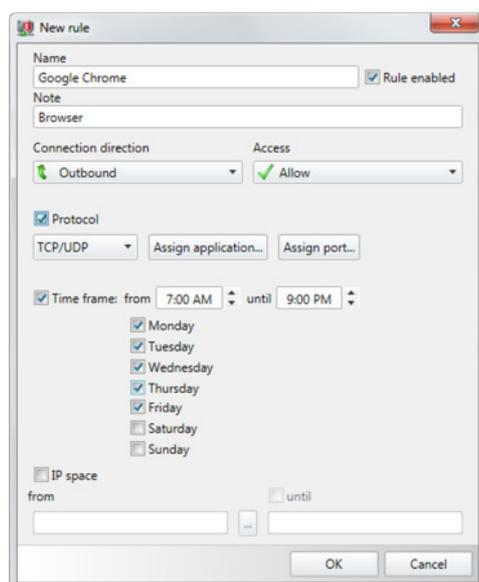


Image 51: G Data Administrator, Firewall, Rule sets, New rule

New rules can be added manually using the NEW button, or using the Rule wizard by clicking the WIZARD button. The wizard can be used to add a rule from the default rule set, copy an existing rule, grant or deny access for a specific application, or open or close a specific port. Manually adding a new rule is most useful when the wizard does not allow enough flexibility, for example when rules should be valid for a specific time frame only. Every rule consists of several components. The basic information that is required is the connection direction and access type. Firewall rules can be set for inbound, outbound, or inbound/outbound connections. Some applications could be allowed to send packets only, while for others, network access can be blocked completely. In addition to the basic information, one or more from the following properties should be selected: PROTOCOL, TIME FRAME or IP SPACE. The protocol dropdown contains several protocols for which traffic can be filtered by the firewall. Additionally, a specific application and/or port can be defined. The APPLICATIONS popup window allows very specific control over the application that should be blocked or allowed. It can be used to simply filter one or more applications' traffic, or to block or allow an application's traffic only if it was started by a specific parent process. The PORTS window lets administrators add single ports or port ranges to the rule. The TIME FRAME option can be used to limit the rule to a specific time of the day or to specific days of the week. This can be useful for both security purposes and company policies. Finally, the IP SPACE option can be used to add an (external) IP or IP range to which traffic should be limited. This can be an IPv4 or IPv6 address, a DNS server, a default gateway server or a WINS server.

Deceptively simple, the `NEW RULE` window can be used to create sophisticated firewall rules. Which rules should be created exactly, depends on the specific network configuration, software inventory, security demands and company policies. Rules from the default rule set can be a great starting point, but almost every rule set will need some fine tuning before it can be deployed to the complete network. Accidentally blocking important traffic can be extremely impractical for end users. While developing a rule set, therefore, it is recommended to test the set on one client before deploying it further. Make sure that communication between the client and its `ManagementServer` is not interrupted, because that will make troubleshooting significantly more difficult.

As a general guideline, it is recommended to configure the firewall to operate in a “whitelist mode”: define rules for known applications only and block all other traffic. This can be achieved by inserting rules at the end of the rule set to block all traffic on all protocols and ports. Every connection attempt that is not covered by one of the other rules will be blocked. Highly secure, this rule set mode can be tricky to configure because it is very hard to predict all types of traffic that a client will generate. A certain amount of time will be required to figure out the normal usage pattern of a client PC and the rules that will have to be configured to cover the intended behavior. One method is designating one client to be a rule set testing ground. This client should have most if not all of the software installed that is in use on the network. Starting with rules from the default rule set, the client can be used to gauge network traffic needs and add the appropriate permissions. The `Firewall log` can be very helpful to get an in-depth view of connection attempts.

Before rolling out any rule sets, it should be made sure that each network zone has its own rule set with the appropriate rules. Rule sets should contain only rules for the communication needs of the specific network zone and its clients. Granting too many applications network access can be a security risk. At the same time, the number of rule sets should be manageable: the more rule sets are defined, the more time will be required to troubleshoot traffic permission problems across the network. For the same reason, rules should be deleted from rule sets if they are no longer being used. For example, when a software product is retired and is no longer used on the network, its corresponding rule(s) should be removed from the firewall rule set. Keeping the rule set clean allows for a better overview over the existing rules and eases troubleshooting.

13.4. End user permissions

A firewall is an essential layer in network security, but at the same time can hinder end user activity on a client machine, if it blocks parts of an essential application. Administrators can move part of the responsibility of firewall management to the end user, by allowing them to enable or disable the firewall, or to change the off-site configuration rule set. This can be practical, because it lets the end user immediately fix permission problems, but dangerous as well: company security policies can be bypassed if the end user is allowed to change rule sets or to disable the firewall completely. It is recommended to proceed with caution and only set these options for problematic client deployments or end users who know what they are doing.

End users can access firewall settings through `G Data Security Client`. A right-click on the tray icon reveals the `FIREWALL` option, which leads to the `G Data Firewall` interface where rule sets can be edited (if the option to do so is enabled and the client is being used off-site). The firewall can be disabled by clicking `DISABLE FIREWALL`. This will disable the firewall without further warning. It will not be re-enabled automatically, so the disabled state persists across restarts. This option should therefore be used with extreme caution: there is hardly any use case in which the end user should be granted such permissions.

13.5. Logs

The firewall component reports blocked applications to G Data Administrator. However, it can be useful to have a more advanced view of inbound and outbound connections. If an application is inexplicably blocked, or if statistics about dropped inbound connections should be obtained, the local firewall log can help. If end users should have the possibility to view detailed logs, the administrator can enable the checkbox Allow the user to change the off-site configuration. This will add a FIREWALL option to G Data Security Client's system tray context menu. Clicking FIREWALL leads to the main interface of the firewall component. The LOG panel shows a detailed overview of all incoming and outgoing connections. The end user can check the connection protocol, initiating application, direction, local port, remote host, remote port and reason for the decision about allowing or blocking the connection.

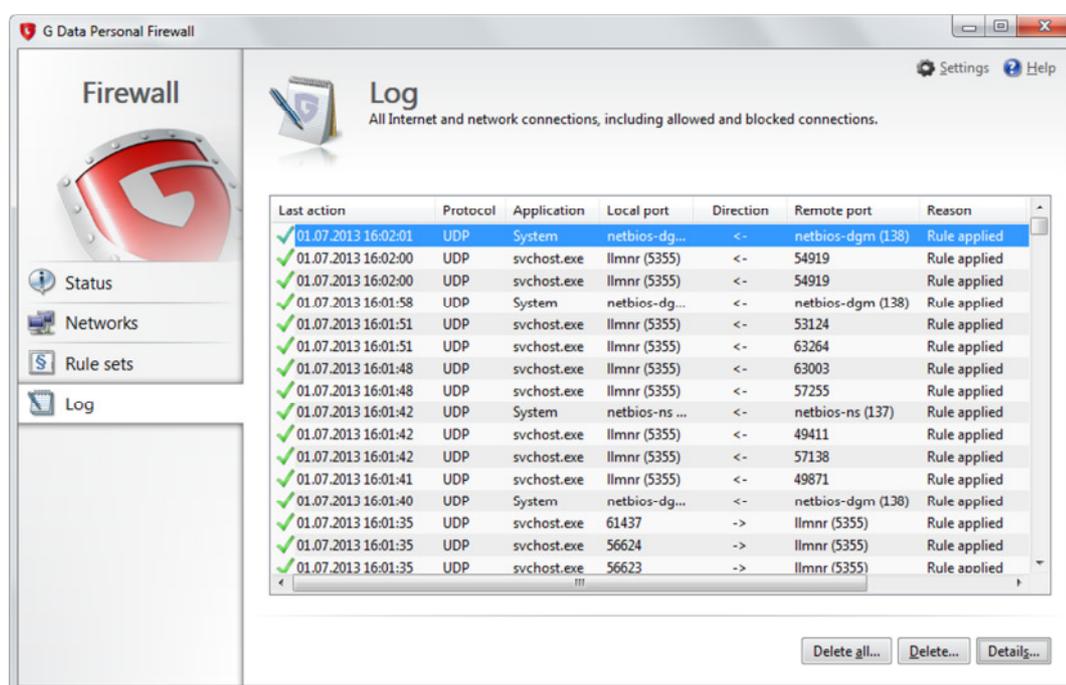


Image 52: G Data Security Client, Firewall, Log

If the logs are not meant to be used by end users, the administrator can choose to directly access the log files themselves. Note that this is an advanced view, which can be overwhelming. The firewall log is saved as a SQLite database file. This file format can be read by a number of database browsers, such as SQLite Database Browser (www.sf.net/projects/sqlitebrowser/). The connection log is stored in the database file LiveStrm.dat, typically located in C:\Program Files (x86)\G DATA\AVKClient\Firewall. The database contains several tables; the connection log is located in the table Connections. Each entry is one connection, and its properties are listed in several columns. The table below lists the most important columns and their descriptions.

Column	Description
AdapterID	Network adapter ID.
AdapterName	Network adapter name.
Allowed	The connection was allowed (-1) or blocked (0).
HasProcess	The process ID that initiated the connection. Can be used to compare against a list of running services.

IPv6	The connection uses the IPv6 protocol (-1) or IPv4 (0).
LocalPort	The local connection port.
Outgoing	The connection was inbound (-1) or outbound (0).
ProcName	The process that initiated the connection. Can be used to compare against a list of running services.
Protocol	The connection protocol ¹² .
Reason	A numerical value representing the reason for the decision about allowing or blocking the connection.
RemoteHost	The remote host.
RemoteIP	The remote IP address.
RemotePort	The remote connection port.
RuleID	The ID of the rule that governed the decision about allowing or blocking the connection. Corresponds to the ID in one of the rule tables of the SQLite database GDFwSvc.dat.

The column `REASON` displays a numeric value for each connection. This value provides additional information about why the connection was allowed or blocked. The value that occurs most often is 7 (`REASON_RULE_MATCH`), meaning that the connection matched a rule in one of the defined rule sets. The complete list of values is as follows:

Value	Reason
0	REASON_FILTER_OFF
1	REASON_FLUSHED
2	REASONS_ASK_USER_CACHE
3	REASON_ASK_USER_WRONG_CHECKSUM
4	REASON_ASK_USER
5	REASON_ASK_USER_NO_FRONTEND
6	REASON_ASK_USER_NO_PROCESS
7	REASON_RULE_MATCH
8	REASON_RULE_MATCH_WRONG_CHECKSUM
9	REASON_SKIP_ID
10	REASON_SUBSEQUENTLY
11	REASON_BY_CONNECTION
12	REASON_ENDPOINT_DOES_NOT_EXIST
13	REASON_ADAPTIVE_MODE
14	REASONS_ANSWER_FROM_OUTGOING
15	REASON_RULESET_DEFAULT
16	REASON_ANSWER_FROM_INCOMING
17	REASON_RULE_MATCH_WRONG_PARENT_CHECKSUM
18	REASON_RULE_MATCH_UNMATCHING_PARENT
19	REASONS_ASK_USER_WRONG_PARENT_CHECKSUM
20	REASON_ASK_USER_UNMATCHING_PARENT
21	REASON_PROCESS_DIED
22	REASON_TCP_ENDPOINT_IS_NOT_LISTENING
23	REASON_RULE_MATCH_WRONG_MODULE_CHECKSUM
24	REASON_JUST_OUTGOING
25	REASON_DHCP_POLICY
27	REASON_FIREWALL_OFF
28	REASON_ICS
29	REASON_WRONG_CHECKSUM_COMMITTED_BY_AV
30	REASON_AUTOPILOT

¹² The protocol number can be looked up in IANA's database of protocol numbers (www.iana.org/assignments/protocol-numbers).

14. PolicyManager

Security layers such as the firewall, file system monitor or on demand protection block infection attempts and provide end users with a safe computing environment without malware. However, many enterprise policies define other types of content that should not be accessible. Inappropriate content is often blocked, applications are blacklisted or internet access is limited. Similarly, the use of external devices with enterprise clients is often limited, by prohibiting the use of USB sticks. All these measures have in common that they are not only configured for network safety, but also to enforce client usage policies. G Data's PolicyManager unifies these issues in an easy-to-use module that can be used to configure clients for exactly the type of use they are meant for.

Before defining any policies, administrators should work out a full overview of the client roles in the network and the desired number of privileges for end users. Clients that are used in the IT department, for example, will have other permissions than sales or R&D. At the same time, there can be enterprise-wide rules that should be applied to all clients: if the company has defined general security policies, they can help in building the corresponding PolicyManager rules. An enterprise-wide USB stick ban, for example, is easy to implement using PolicyManager's Device control panel. Application control, Device control and Web content control can be combined to create rules that ensure that no sensitive data can leave the enterprise network, be it via file sharing applications, USB sticks or cloud storage.

As with all other aspects of security, policies should not be immediately deployed to all clients. PolicyManager can be used to restrict several aspects of client usage, but the ability for end users to effectively use the PC should not be limited. To properly gauge its effects, every policy change should be tested on one or more test clients. Only if it functions as intended, it should be added to the default policy for all clients.

The PolicyManager module is divided into four panels: APPLICATION CONTROL, DEVICE CONTROL, WEB CONTENT CONTROL and INTERNET USAGE TIME. Each of them controls a specific aspect of client use and can be enabled or disabled on its own. As usual, settings apply to the clients that are selected in the client management area. PolicyManager modules can be enabled for normal client end users, or for normal client end users and administrators. The latter setting is recommended for optimal security, specifically if local administrators should not be able to circumvent the policy rules.

14.1. Applications

In security as well as policy, allowing end users to run any applications on a client can be a source of major headaches. Software from an unknown source can contain malware. Furthermore, a large group of applications could be designated as unwanted software. In most enterprise scenarios, for example, end users do not need peer-to-peer download software. Instant messaging applications are another type of software that is often blocked. Application control allows administrators to manage and enforce application rules using a blacklist or whitelist, configurable per client.

Before any decisions about application whitelisting or blacklisting can be made, administrators should be aware of the software that is running on network clients. The list of software does not only include packages that have been officially deployed. If there are no restrictions in place as to the installation of new software, end users may have installed any applications. To gain an overview on the software that is in use on a client, use the CLIENT module. Its SOFTWARE panel shows all programs that have been installed on a client. This information can be used to decide which applications are allowed and which ones not. Software can be preliminarily sorted by using the software inventory's

blacklist and whitelist functions. For each program, decide whether it should be allowed on client PCs or not. Right-click the entry and choose PUT ON WHITELIST or PUT ON BLACKLIST. The buttons GLOBAL BLACKLIST and GLOBAL WHITELIST show the respective lists, which can then be used as basis for decision making in the APPLICATION CONTROL module. Pay attention to the VENDOR column, as the vendor name can be used to create a blacklist or whitelist rule in PolicyManager.

In order for Application control to function correctly, the file system monitor component of G Data Security Client needs to be enabled, as its continuous file system checks help recognize application starts (see chapter 8.2). Per client, the choice can be made whether Application control should function in blacklist or whitelist mode. In the blacklist mode, all applications that are defined in the application list will be blocked when an end user tries to run them. Whitelist mode only allows the applications on the list to be run and blocks all others. Even though it is possible to define a different mode for different clients, it is recommended to use the same mode for all clients to make management easier. Whitelist mode is the most secure mode: only applications that are deemed safe can be run. However, defining a whitelist requires quite a lot of tests, as administrators should find out about every program that should be allowed, to avoid end users encountering blockades while trying to work. Blacklist mode is the alternative: administrators will only have to define programs that they do not want to allow. The downside of a blacklist is that it does not automatically take care of programs that were not known to the administrator at the time the blacklist was defined. If an end user installs a program by themselves, it is not automatically blocked.

To add a rule to a list, click NEW to add a custom rule or DEFAULT RULES to choose one from a list of defaults. The default rules cover a host of popular third party programs. Adding a custom rule is not complicated and can either be defined by vendor, file or directory. A vendor rule blocks or allows executable files based on their vendor string. Especially when using whitelist mode, make sure that at least the operating system and G Data security components can be loaded properly by adding a vendor rule with the string *Microsoft** or *G Data**, or by using the respective default rules. If a vendor should be generally blocked, with the exception of one or more programs, their executables or folders can be defined as an exception to that specific rule. File rules can be added by entering properties of the file to be blocked, such as file name, MD5 checksum, product name, file version or copyright. These fields can be entered manually, or populated automatically by using the DETERMINE FILE ATTRIBUTES option. This will allow an administrator to select the actual file to make sure the properties are entered correctly. The beginning or end of any of the property strings can be followed by an asterisk. This is especially useful when blocking only a certain version range. A directory rule lets administrators select any folder from which executables should be blocked or allowed, optionally including subfolders.

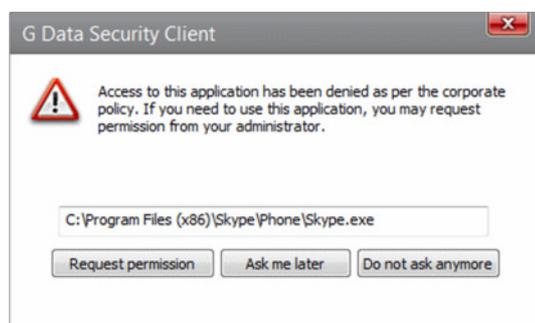


Image 53: G Data Security Client, Application control

When all rules have been defined and Application control has been enabled for the appropriate clients, it is recommended to test if all rules work correctly. When an end user tries to start a blocked

application, G Data Security Client will automatically prevent access. If the administrator enabled the option **THE USER CAN REPORT BLOCKED APPLICATIONS**, Security Client will present a popup with application details. The button **REQUEST PERMISSION** will add a report to the **REPORTS** module of G Data Administrator. Using this report, the administrator can directly add a new rule to Application control, whitelisting the app if the end user should be able to use it. Leaving this option disabled is recommended for clients where applications should be blocked without user interaction or feedback.

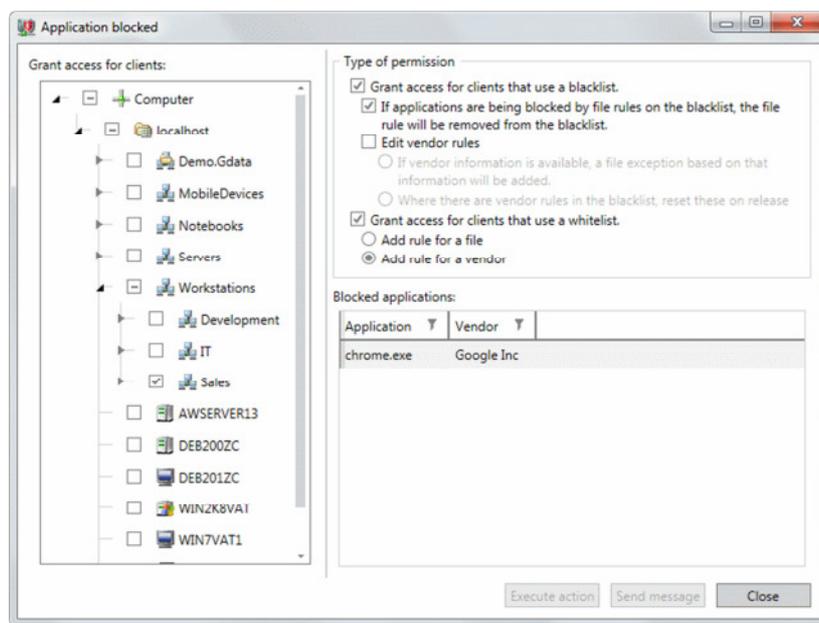


Image 54: G Data Administrator, Reports, Application blocked

The report that Application control adds to the **REPORTS** module allows for some flexibility in defining a rule. The application that was blocked is listed in the bottom right corner. On the left, the client(s) for which the new rule should be made can be selected. By default, this is the client from which the report was sent, but applications can be white- or blacklisted for any of the network clients. The **TYPE OF PERMISSION** can be tweaked both for blacklist and whitelist modes. For clients that are using a blacklist, the rule that blocks the affected application can be removed. If it is being blocked by a vendor rule, the file can be added as an exception to that rule, or the rule as a whole can be removed. For clients that use a whitelist, the program can be added as a file rule or a vendor rule.

14.2. Devices

Most malicious threats originate from the internet, but removable devices remain a very popular attack vector for malware. Removable storage media like USB sticks and CD-ROMs can contain viruses, but administrators should not only worry about attacks: sensitive files could also be leaving the enterprise network through removable media, for example when an employee copies an important database to a USB stick. Another delicate device category is webcams, which are often built into laptops but can be disabled for privacy reasons. The **DEVICE CONTROL** panel allows administrators control over these device categories across all network clients.

Per device category, the administrator can define access rights. These are the same for floppy disk, CD/DVD, removable storage media and web cam: **READ/WRITE**, **READ**, or **DENY ACCESS**. By selecting **Read/Write**, users are allowed full access to the selected device(s). **Read** access is a useful option for situations where end users need to be able to copy data onto a client, but should not be allowed to copy data from it. **Deny access** is the best option if a device needs to be blocked completely. This is

the most secure feature: devices cannot infect clients with malware, and they cannot be used to transport data out of the enterprise network.

Blocking device categories completely is the easiest part of using Device control. However, there can be circumstances where end users or administrators want to use a device locally, in spite of a network-wide ban. Rather than enabling READ or READ/WRITE access for all devices in a category, the whitelist feature lets administrators define specific devices that can be used. If a device category has been set to READ or DENY ACCESS, the whitelist can be used to add READ or READ/WRITE permissions to a specific device or medium. Before rolling out a device policy, essential device or medium permissions should be whitelisted, to make sure that no workflows are interrupted.

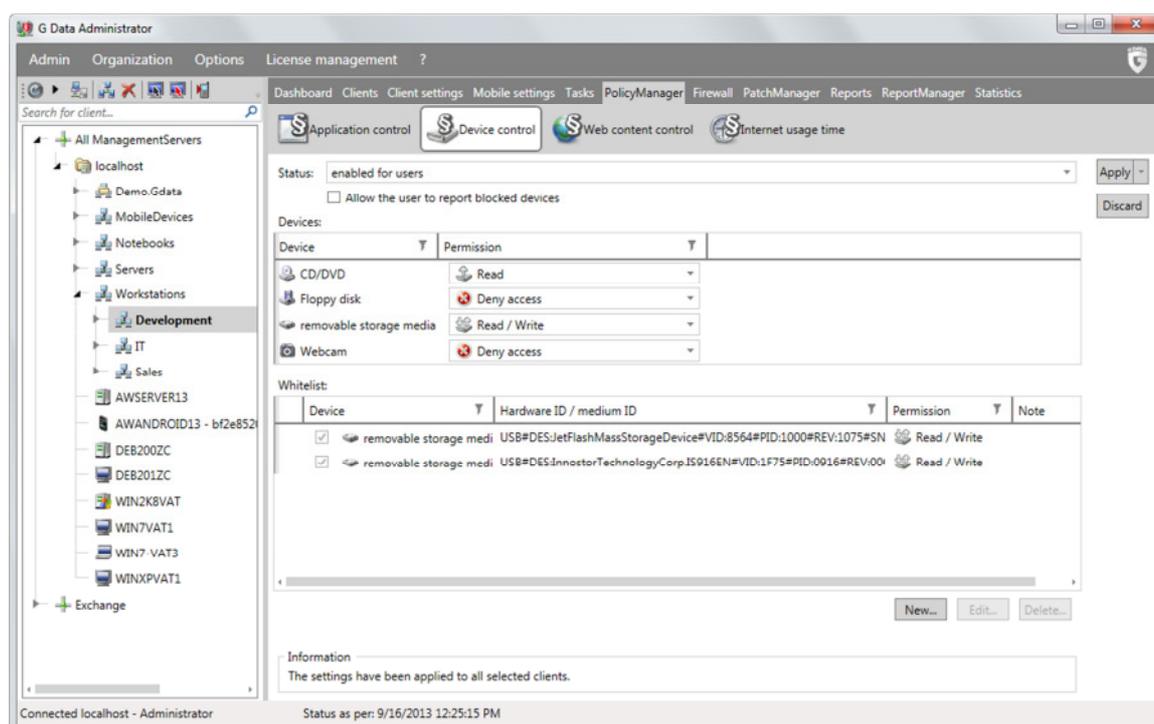


Image 55: G Data Administrator, PolicyManager, Device control

Clicking the NEW button brings up the whitelist entry popup. The window lists only those device categories that have been restricted in the device category list; device categories that have been assigned READ/WRITE permissions are not being blocked and hence do not need whitelist entries. The field HARDWARE ID/MEDIUM ID is used to define the device or medium that should be added to the whitelist. Using the ... button, the correct ID to be whitelisted can be determined. The popup window allows a search across local devices or devices on any of the network clients. By selecting a hardware ID, a complete device can be whitelisted. For example, a system where all DVD drives are blocked, can have one DVD drive defined as an exception. A medium ID, on the other hand, helps whitelist a specific medium. For example, on a system with blocked DVD drives, permission can be added to use a specific CD or DVD. This is also useful for clients that are not allowed to use any removable media, but need access to one specific (company-managed) USB stick.

Similar to Application control, the Device control module can do its work with or without user interaction. Select THE USER CAN REPORT BLOCKED DEVICES to provide the user with the possibility to request access whenever a device is blocked. The popup window will offer an option to request permission to use the device or medium. This will generate a report in the REPORTS module. Based on the report, a whitelist exception can be added. The device category itself can be enabled for selected clients, or

only a specific device or medium based on device ID or medium ID. The possibilities for whitelisting depend on the data that are available in the report. Be careful when adding device exceptions. The decision should not be based on user requests alone; if possible, obtain the medium that the exception was requested for and verify that it is indeed necessary. When adding an exception, work with medium exceptions instead of hardware exceptions wherever possible. This will limit the risk of unauthorized device use and malware infections.

14.3. Web content

A popular part of enterprise policies is limiting access to specific websites. There can be many reasons for blocking sites. Productivity is one: end users can be prevented from visiting sites that have nothing to do with the tasks they are working on, in order to block popular time sinks (such as online games or social networks). Additionally, sites can contain illegal or inappropriate content. File sharing sites, adult content, or sites propagating hacking are in almost all cases irrelevant to work and should be blocked. This type of sites is often found in the shadier parts of the internet, where malicious web hosts or hacked ad networks could try to infect visitors with malware.

Before blocking any of the categories included in PolicyManager's WEB CONTENT CONTROL, an inventory of essential sites should be made. To ensure that important sites are not accidentally blocked, they should be added to the GLOBAL WHITELIST before deploying any policy. If there are any sites that need to be blocked by name, they can be added to the GLOBAL BLACKLIST.

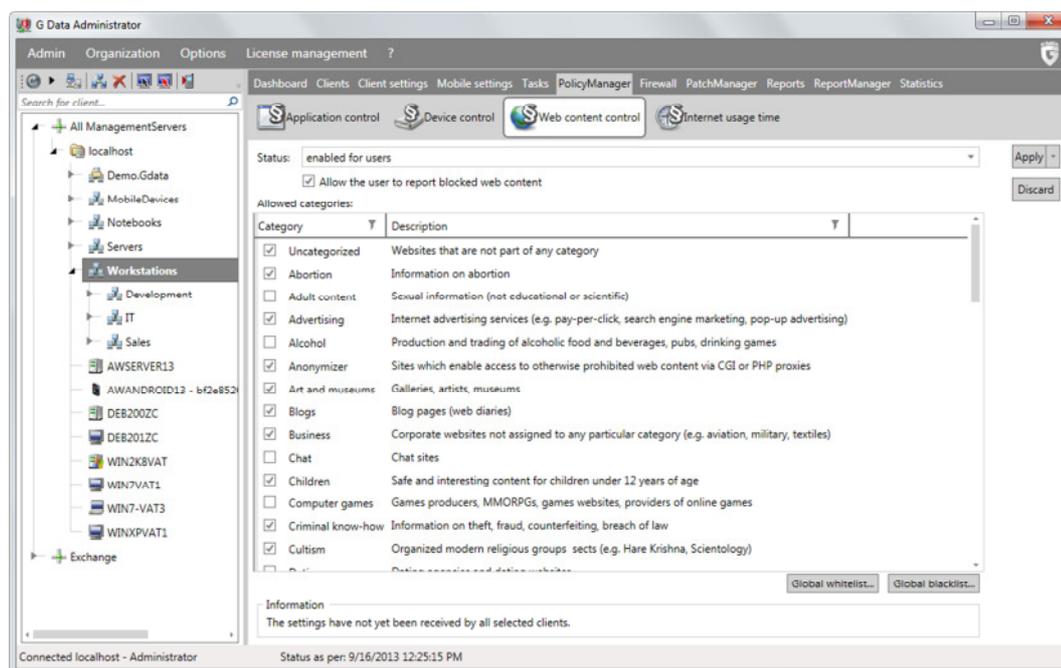


Image 56: G Data Administrator, PolicyManager, Web content control

To save time and effort, administrators can block predefined website categories. Several categories are available, grouped by type of content. Each category consists of a list of blacklisted URLs that will be blocked. For each network zone, administrators should decide which websites categories should be blocked. It is recommended to block illegal content at the very least (depending on local legislation, this includes categories like Criminal know-how, Drugs, File sharing, Gambling, and Hacking). Depending on network policies and administrator wishes, other categories can be enabled (for productivity, categories such as Blogs, Chat and Social networks; inappropriate content can be

blocked using categories like Adult content, Hate, or Nudity).

With the correct client or group selected in the client management area, the categories to be blocked can be selected. If the checkbox in front of the category name is selected, access to websites in that category is allowed. To block access, deselect the category. Any combination of categories is possible. Since categories are provided for a very wide-ranging set of scenarios, blocking access across all of them at the same time is not recommended, since this will severely impede website access. Administrators should make a specific selection based on the types of sites that employees should and should not be able to visit.

In addition to the category list, sites can be blocked or allowed by using the GLOBAL WHITELIST and GLOBAL BLACKLIST. The global whitelist can be used to define sites that should be allowed to be visited, regardless of category selections. Conversely, the global blacklist will block websites, even if they would not have been blocked based on the categories that are in effect. The global whitelist and blacklist are in effect for the whole network to allow administrators to quickly define exceptions for all clients. For example, websites that are essential for workflow processes can be added to the whitelist, or the company's own website.

Web content control relies on the HTTP traffic scan layer (see chapter 8.1). This means that the option to process internet traffic should be enabled for every client on which web content control is to be deployed. PROCESS INTERNET TRAFFIC (HTTP) can be enabled on the WEB/IM tab of the CLIENT SETTINGS module. If the client uses a proxy server, it can be enabled on the same tab. When the end user requests a URL in the browser, G Data Security Client checks the URL against G Data's central URL list to requests its category. If the category is blocked, or the site is on the global blacklist, the request is denied and the page is not loaded. If there are latency issues, and the category lookup delivers no result within 1000 milliseconds, the site will be loaded. If THE USER CAN REPORT BLOCKED WEB CONTENT has been selected, a popup window will be shown. The user can request access to the website, which will generate a report in the REPORTS module. The administrator can then either allow access to the complete category (for any of the clients or groups), or add the website to the global whitelist.

14.4. Internet usage time

The fourth and final PolicyManager module focuses not on filtering internet content, but completely blocking it. As with all policies, there is both a policy and a security component to it. Policy-wise, administrators can choose to restrict internet access to specific times of the day, to restrict it to a maximum cumulative amount of time, or to completely block it. Employees that do not need internet access can be forced to focus on their tasks, or be allowed to use the internet only during a lunch break. Moreover, internet access restriction provides security by reducing the amount of time that the attack vector is available. However, it needs to be carefully considered in which scenarios it can be used. Blocking internet access severely restricts workflows that depend on obtaining or publishing information online.

The Internet usage time panel of PolicyManager shows two important components. On the left, the grid shows at which times on which days of the week internet access should be granted or restricted. Using the time grid, several configurations can be enacted. Click and drag, or use the Ctrl + Click or Shift + Click combinations, to select several time slots at once. Right click to choose between blocking and allowing internet usage for those time slots. For example, to make sure that internet access is only possible during lunch break, select all time slots (Ctrl + A), right click and select Block time. Then select lunch break hours, right click and select Allow time. If an end user tries to access a website when internet access is blocked, a warning page will be displayed in the

browser. Note that internet access restrictions could be circumvented by changing the local time setting. If the setting has not been enabled already, it is recommended to prevent end users from changing local time settings using a group policy object.

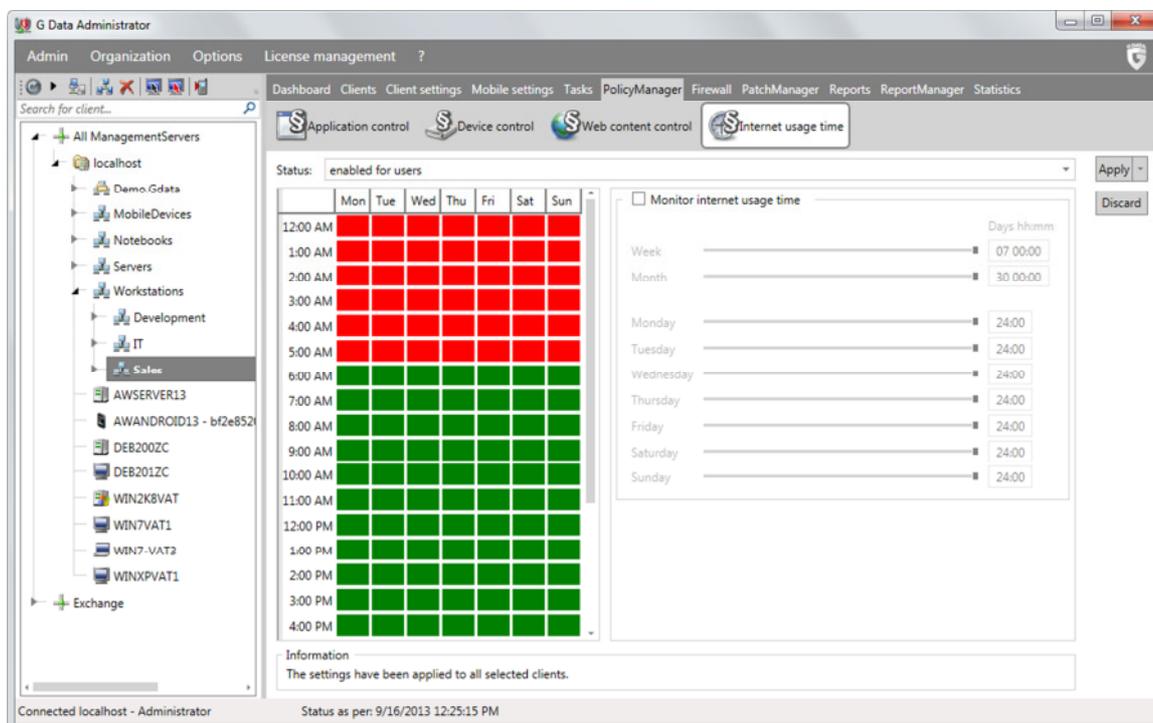


Image 57: G Data Administrator, PolicyManager, Internet usage time

On the right, a set of sliders can be used to define daily, weekly or monthly usage limits. These limits are enforced separately and in addition to the time grid. The sliders can be enabled by ticking MONITOR INTERNET USAGE TIME and indicate the maximum amount of time internet can be used. By default, internet access is available all the time: 30 days per month, 7 days per week and 24 hours a day. By dragging the sliders, this amount can be reduced. Alternatively, time can be entered manually. For example, entering 04 20:05 in the Week field corresponds to 4 days, 20 hours and 5 minutes of internet usage allowance. Once the internet allowance for a certain period of time has been exceeded, users will see a warning page if they try to open a website. If there are any conflicting time entries, the smallest amount of time will be used. For example, if a time limit of four days per month is defined, but a weekly limit of five days, the software will automatically limit Internet usage to four days.

15. Patch management

Patches often repair security vulnerabilities through which attackers may gain access to systems running the affected software. In responding to security emergencies, rapid deployment of patches is important. A complicating factor, the release of a patch actually stimulates hackers to develop and exploit the security bug, due to the public release of information about the vulnerability that typically accompanies patch releases. By reverse-engineering patch files, attackers can obtain the information necessary to stage an effective attack. This puts extra pressure on administrators to timely patch their systems. Patch management helps speed up patch deployment and improves the efficiency of the complete process by coordinating and standardizing patch deployment procedures.

Not applying patches leaves systems exposed. Attackers who take advantage of software bugs can, depending on the severity of the security hole, gain access to files saved on a PC, execute programs, take over other PCs in the corporate network, or worse. While a malware infection through a drive-by download or an attack from hackers are annoying for home users, corporate networks are especially vulnerable. The stakes are much higher: the mere presence of a hacker in the company network compromises data integrity, and can lead to the loss of data if one or more systems are irreparably damaged. Further threats include downtime for critical systems, intellectual property theft, loss of reputation, or excessive costs for legal defense if customer data is lost.

Standardized patching procedures help prevent successful exploitation of software bugs by hackers. However, patch management is not the only measure that should be taken. Even if software is fully patched, hackers may be aware of bugs that have not been discovered yet by the software vendors. Business networks and clients should always be protected by security products on a client level, providing measures such as signature based malware recognition, heuristic scanning and file reputation management, as well as protection on the network level.

15.1. Patch management cycle

Patch management is important for each computer, be it a home device or corporate workstation. In any case, the availability of new security patches should be actively monitored and fixes should be deployed as soon as possible. However, the way in which patches are managed and deployed is mostly a question of scale. For home users, Microsoft's built-in Windows Update can provide security fixes for Windows in a completely automated way, and more vendors are moving towards a fully transparent, automatic updating process (Adobe with its Adobe Reader and Adobe Flash Player software, as well as browsers like Mozilla Firefox and Google Chrome). They are regularly alerted by built-in updaters of other software. Tech-savvy users may choose to use commercial update notification software, to make sure no patches are being skipped¹³. Nevertheless, home PCs are among the most insecure devices in regards to patch management, mostly due to indifference or a lacking sense of urgency on the side of the end user. The complexity of keeping an overview of installed software, its vulnerabilities, and its patches is already overwhelming for single PC users – let alone administrators of business networks with any number from five up to thousands of clients. This is where a standardized, recurring patch management procedure can help, by reducing the time required to take an inventory of software and vulnerabilities, and by automating the deployment. An effective patch management procedure clarifies which responsibilities lay with whom, tracks all changes that are being made, provides a rollback method, tests all proposed

¹³ Several commercial parties offer update notification software, for home users as well as small businesses. Some of these solutions only function as reminders, others almost resemble patch management solutions. See, for example, UpdateStar at www.updatestar.com.

changes extensively, and announces changes to all involved parties.



Image 58: Patch management cycle

The patch management cycle can be broken down into different stages. Note that it is a cycle, not an event-driven process. Patches should be proactively deployed, therefore patch management should be proactively carried out. Each step of the cyclical procedure should be explicitly defined and assigned to someone. Depending on wishes and requirements, companies can merge stages by bundling them and assigning them to the same person, or define further action points as required. Existing change management and release management standards can be (partially) integrated. Some steps of the procedure can be automated, most notably the deployment, but several key actions will have to be carried out manually for each cycle. To optimize this process, planning is crucial¹⁴.

The rate of recurrence of the patch management cycle depends on the resources that can be made available to execute it, as well as on the rate with which patches are published for software that is in use in the network. A major source of patches is Microsoft, whose Windows operating system is served with the latest security updates and improvements monthly. Some other vendors have chosen to align their update cycle with Microsoft, most notably Adobe (whose Adobe Reader and Flash Player software are along the most widely deployed installations)¹⁵. These predictable patch cycles are scheduled to occur on a fixed date, in the case of Microsoft and Adobe every second Tuesday of the month. This allows administrators to plan a recurring, reliable procedure to roll out the latest updates every month. Other vendors release updates less frequently, such as Oracle, which has chosen to distribute security patch bundles for the widely used Java platform every three to four months¹⁶. Some vendors choose to release patches whenever a critical issue comes up. For patch management planning, this quick response time comes at a price: it may take some time before an unpredictably released patch is thoroughly tested and deployed. On the other hand, rigid patch bundling can cause some security issues to go unpatched for a significant time until the next patch release date.

¹⁴ Chapter 15.2 will discuss how to carry out the various steps of the patch management cycle using G Data PatchManager. For a more theoretical overview of the different steps, see G Data TechPaper #0271 – Patch Management Best Practices.

¹⁵ Adobe releases quarterly and out-of-band updates for its Acrobat and Reader products (helpx.adobe.com/acrobat/release-note/release-notes-acrobat-reader.html). Its Flash Player has recently been moved to a rapid release cycle (blogs.adobe.com/flashplayer/2012/09/flash-runtime-rapid-release-cycle.html).

¹⁶ See www.oracle.com/technetwork/topics/security/alerts-086861.html.

15.2. Standardized patching policy

Before planning the monthly steps for a patch management cycle and assigning responsibilities, some standards have to be defined. A patch management policy helps decision making during the cycle. The policy should cover questions about patching strategy. Should all available patches be installed by default, or will there be a classification, possibly based on the severity of the security holes they remedy? Will patches be installed proactively (to plug possible security holes) or reactively (only when problems arise), or a combination of both? To prevent spending unnecessary time on patch-by-patch decisions, it is recommended to set as many generalized rules as possible. At the same time, simply installing every available patch is not a solution: to prevent network and system load and compatibility problems, conscious choices need to be made.

Other parameters that the patch management cycle depends on include installation standards, network standards, and application security configuration. While each cycle includes making an inventory of the current state of the network, effort is greatly reduced when the network environment has been set up following a standard. Define which software installations are allowed, and which machines are provided with which software. Using whitelists or blacklists, the network software inventory can be limited, greatly helping patch deployment. The same goes for security configuration, user accounts and passwords.

By standardizing policies across the network, there will be fewer exceptions to be handled during patch deployment. However, it can still be fruitful to define actions to undertake in case exceptions do occur. Administrators should not be surprised by machines that turn out to be unreachable during deployment, patches that cause compatibility problems during testing or deployment, or security issues that do not seem to have been mitigated successfully after patching. Swift redeployment, software reconfiguration and incident escalation should be defined as part of the patch management policy.

An important aspect to keep in mind is the deployment of new machines in a corporate network. While outside the strict scope of patch management, it is important to make sure that system images with which new clients are deployed are maintained with the latest patches. If outdated systems get deployed, there is a vulnerability window that will only be closed once the newly deployed machine has become part of the patch management cycle and its inventory has been listed.

15.3. PatchManager

Patch management is available as an optional module for all users of Business and Enterprise products. Combined with the existing modules of G Data AntiVirus, ClientSecurity and EndpointProtection, functional support for all phases of the patch management cycle (see chapter 15.1) is provided.

15.3.1. Step 1: Inventory update

Firstly, it is important to take and keep an inventory of machines in the network and their software and hardware. To support the patch management cycle, it should be known which software versions are in use on the company network at any time. The CLIENTS module lets administrators access a full list of installed software for each network client. The inventory can be organized to provide different types of information. The default view shows a flat list of all software that is installed on the selected clients. The listing includes the installation date, the software vendor and the currently installed

version. By grouping the items according to vendor and name, for each product a quick overview is available to check if the latest version has been installed on all machines.

At this stage, it is recommended to check if network clients are running any software that is not part of the standard deployment. Administrators cannot be aware of potential security risks for all software in existence. Using the Software inventory helps spot unsanctioned software installations. Administrators can decide to either add the software to their official deployment list (Whitelist), or to remove them and block them from being installed or run (Blacklist). Users of G Data EndpointProtection can use the PolicyManager module to apply network-wide policies, whitelisting or blacklisting software to control deployment.

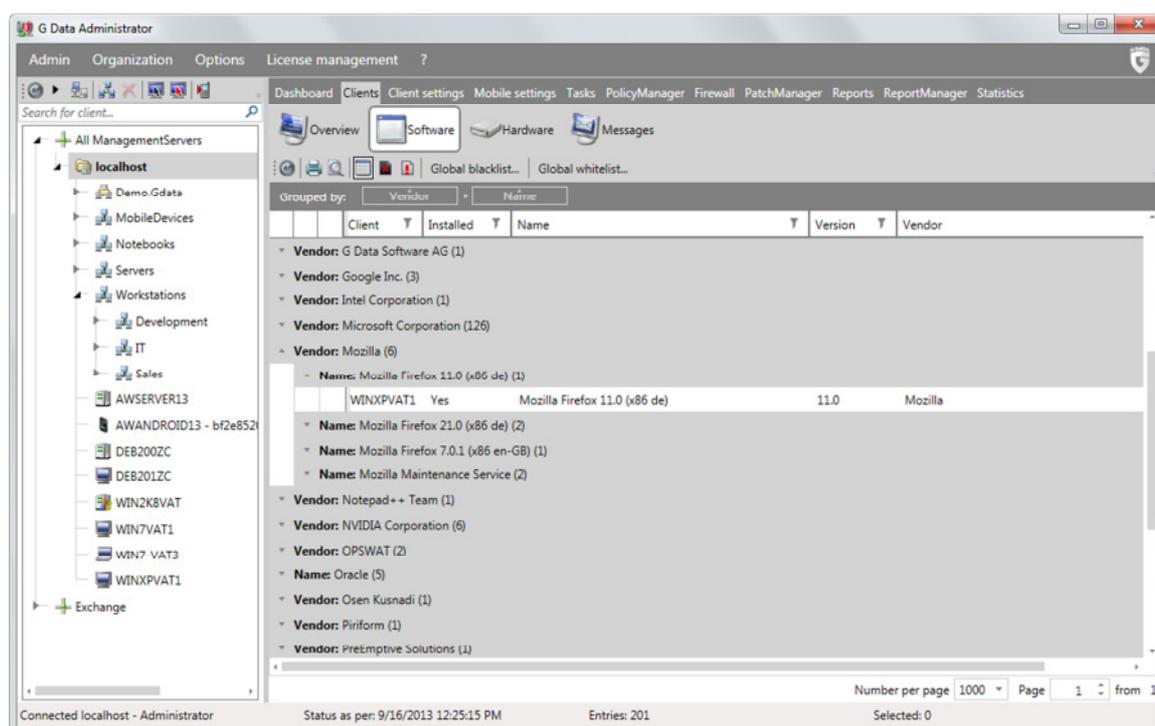


Image 59: G Data Administrator, Client Software Inventory view

Not only is it important to keep track of software; successful deployment also depends on physical prerequisites like network load and hardware specifications. The latter can be listed using the HARDWARE inventory function. A wide range of specifications can be tracked. Physical specs, like CPU speed and the amount of internal memory, help predict patch deployment speed and performance. Important is knowledge of the amount of free disk space available, to prevent patch deployment from generating errors. Additionally, bios and motherboard firmware versions can be tracked, to compare against newly published firmware.

15.3.2. Step 2: Information gathering

As soon as an inventory has been established, administrators should keep up with information about the latest patches to compare to their existing stock. The PatchManager module provides a list of all available patches for a wide range of products on the PATCH CONFIGURATION tab. The database is updated automatically as soon as vendors publish a new patch. By default, the list is grouped by VENDOR, PRIORITY, and PRODUCT. This allows administrators to quickly look up patches for a specific product. More information, and often full release notes, can be obtained by right-clicking on a patch and checking its properties.

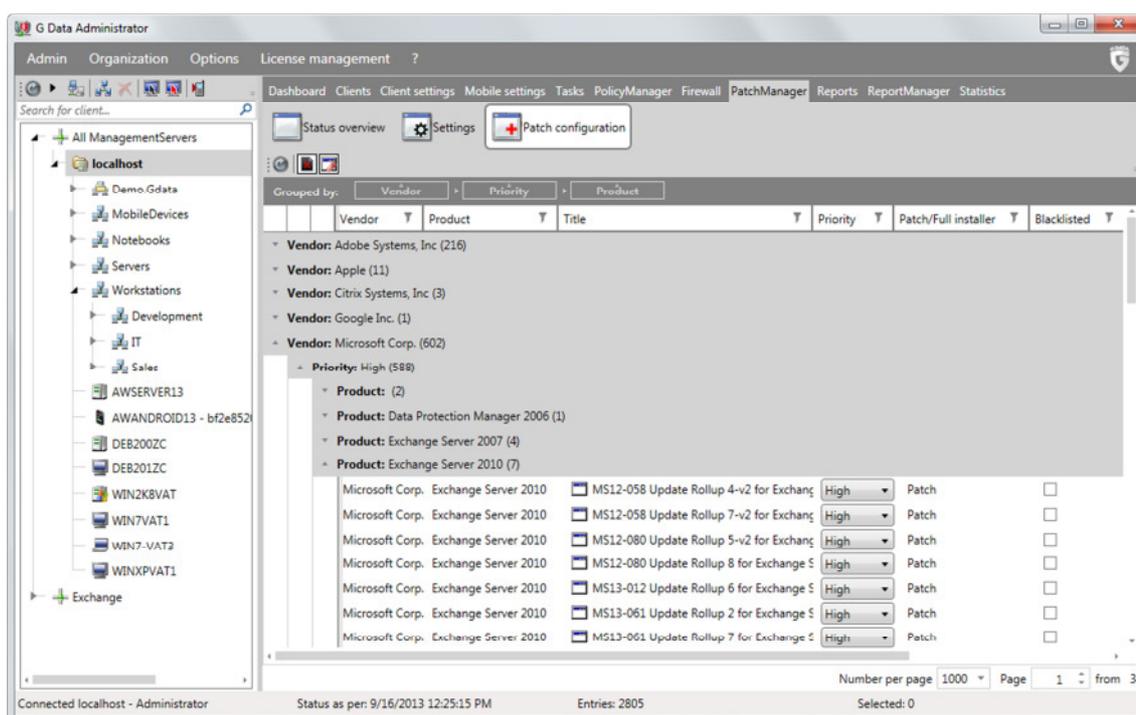


Image 60: G Data Administrator, PatchManager Patches view

15.3.3. Step 3: Strategy and planning

Administrators can directly check one or more patches against their client systems, if it applies to a product that is in use in the network. To check a patch for applicability, right-click it and click CHECK PATCHES FOR APPLICABILITY. This schedules a PATCH APPLICABILITY JOB for the specified client(s). Alternatively, an automatic scan can be carried out for each new patch that is added to the database. Using the TASKS module, plan a PATCH APPLICABILITY JOB that is executed as soon as a new patch is available. PatchManager then checks the new patches for applicability across all specified clients.

PatchManager does not automatically install patches, to allow administrators to select patches and plan patch tests beforehand. After scanning for applicability, select the appropriate server or client(s) in the client management area and open the STATUS OVERVIEW tab of PatchManager. By default, the list is grouped by STATUS, PRIORITY, VENDOR and PRODUCT. This helps to quickly locate patches that are applicable, not applicable or have already been installed. Patches that are applicable for the client system(s) are the ones that need to be reviewed, tested and finally deployed.

To help decide whether to deploy a certain patch or not, PatchManager provides a set of information for each patch. In its list overview, the PatchManager module shows the products that a patch applies to, as well as its release date, its official title, and its priority. For each patch, a full description and usually a URL to the official release notes are provided. These pieces of information help administrators decide how severe a certain vulnerability is, and how quickly its patch needs to be deployed. Patches with a higher severity should be installed with a higher priority than non-critical patches. At the same time, not all software that is in use across the network is as important: critical applications should always be patched before less critical applications. This is where the patch management policy is relevant, by helping to decide quickly which patches to deploy in which order (see chapter 15.2).

The important point to remember is that not all patches should be installed by default. The point of

patch management automation is not to take decision making out of the equation, but to provide enough details to make informed decisions, and to streamline the deployment process. PatchManager provides as much information as it can, but the decision to test and finally deploy a patch, is always up to the administrator.

15.3.4. Step 4: Testing

Once it has been decided that a specific patch will be deployed, the testing procedure can start. It is recommended to use a set of representative machines to test patches. These machines should be similar to the clients that are actually in use, in order to test for possible problems without disrupting the actual clients. However, not every administrator will have access to enough machines to build a small-sized replica of their network. Virtualization is the recommended method; if there is really no other solution, a non-vital subset of the network can be used. In this case, using G Data Administrator, a test environment can be organized in one or more groups. Patches can be deployed to one or more clients in one or more groups, to observe the installation and its effects.

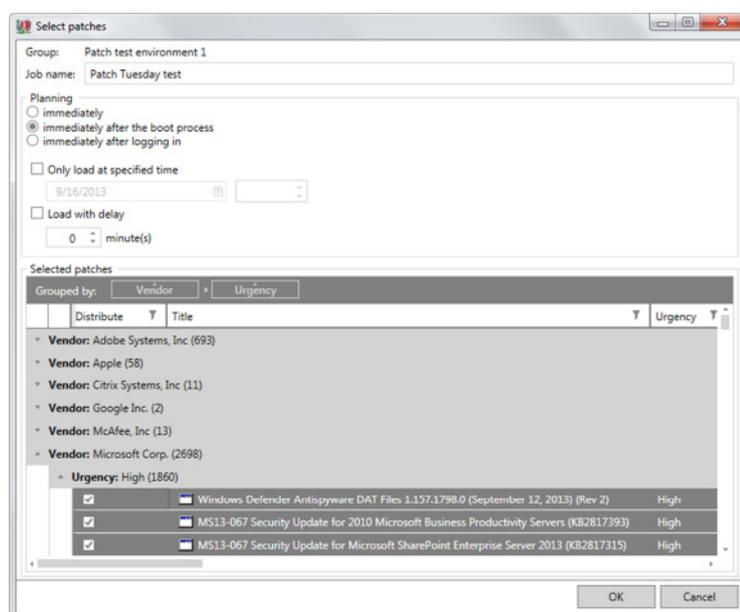


Image 61: G Data Administrator, Tasks, Software distribution job (Test)

Ideally, the test environment features all possible client configurations that exist in the network. Errors and exceptions can only be located by emulating the deployment experience as closely as possible to how it will be carried out in the real network. This means setting up several physical client PCs, each with a representative configuration of one of the network client roles. Complications arise when planning patch tests for machines that have a central role in the network, such as servers. It is very hard to set up a proper testing environment that takes into account all aspects of the network configuration. One option is the deployment of a virtualized testing environment, with virtual server(s) and clients. While this does allow for software and configuration problems to be located, the physical aspect of patch deployment (e.g. bandwidth, free disk space) cannot be tested properly on a virtual machine. Alternatively, a non-vital part of the actual network can be designated as testing grounds. This allows an administrator to test patches in a very realistic environment. As long as patches are only test-deployed to non-critical machines, an actual network test can be very helpful. However, besides the obvious increased risk of complications, this approach blends the testing step with the deployment step. Keeping track of machines used for tests and separating

them from the rest of the clients can become complicated very quickly.

To deploy one or more patches to a test group, select the group in the client management area. Open the **TASKS** module and create a new **SOFTWARE DISTRIBUTION JOB**. Select the patch(es) to be distributed and specify at which time this should occur. Selecting the patch can be made easier by grouping the patch list by **VENDOR** or **PRODUCT**. Repeat this process with all appropriate patches and for all appropriate test groups. It is recommended to test only one patch per system at the same time, to be able to pinpoint possible problems on a specific patch.

During the testing period, as well as the verification phase after deployment, the **REPORTMANAGER** module can assist in finding out what the status is of patches being deployed, and which machines are potentially generating errors (see chapter 6.3). ReportManager lets the administrator select several modules to be combined into one report. Its **PATCHMANAGER** category provides several useful options, such as the patches most frequently not installed or computers with unexecuted software distribution jobs (which may point to installation problems), or computers with the most frequent patch requests or refusals (for analysis afterwards).

In addition to the **REPORTMANAGER** module, patch testing status can also be located in the **TASKS** module itself. Open the relevant task and check the details to see the status for each patch. If it appears that a patch has not been deployed successfully, update the **SOFTWARE** inventory for that client to double-check. If the patch cannot be deployed, check the system locally and try a manual patch deployment. If a patch is causing problems during the testing phase, it should never be deployed automatically on a large scale.

15.3.5. Step 5: Schedule and assessment

After finishing the testing stage, the actual deployment can be planned. With all applicable patches located and tested, a schedule can be set up. Using the patch management policy, decide in which order the patches should be deployed and to which (groups of) machines at first. Use the **MESSAGES** function of the **CLIENTS** module to notify clients of the patch schedule and to warn them for eventual reboots.

15.3.6. Step 6: Patch deployment

For patches that have been properly tested, a **SOFTWARE DISTRIBUTION JOB** can be planned. Use the **TASKS** module to schedule a **SOFTWARE DISTRIBUTION JOB** with the appropriate patches for the appropriate clients. To prevent interference with end user workflows, patches can be scheduled to be deployed at a specific time, or directly after the next boot or login. An optional delay prevents patches from being deployed while other system-intensive processes may be running.

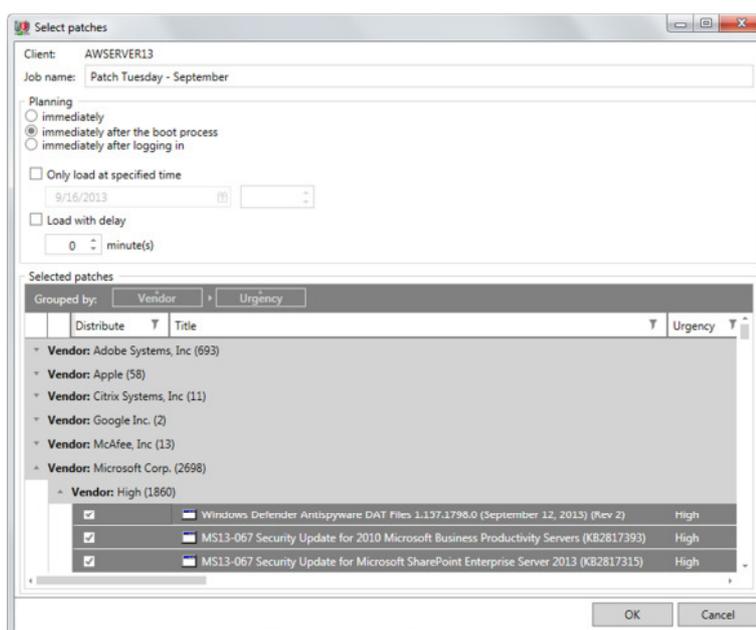


Image 62: G Data Administrator, Tasks, Software distribution job (Deployment)

15.3.7. Step 7: Verification and reporting

To verify and evaluate patch deployment, the inventory tools can be of great assistance. Additionally, the PatchManager module offers the possibility for direct user feedback. If the administrator enables the respective option, end users can request patches to be rolled back, due to performance or compatibility issues. Patches that are applicable to the system, but have not been deployed yet or will not be deployed at all, can be requested by end users in case there is an urgent need to patch a product. The distribution and rollback request system integrates directly with the PatchManager module and allows the administrator to plan a distribution or rollback job directly from the REPORTS module. Consider the following example: an end user can no longer use application 1 and is waiting for a patch to be deployed. During the testing phase, the administrator discovers compatibility issues with application 2, and decides the patch will not be deployed to the network. The end user does not use the affected feature of application 2, and decides that the patch for application 1 should be deployed anyway. Through bidirectional patch distribution the user can request the patch to be installed. Upon approval of the request, patch distribution will take place like it would have done in a normal deployment scenario.

16. Mail server security

To keep malware from infecting end user hardware, such as desktop computers or smartphones, protecting only these endpoints will not do. Other network components should be protected to make sure that threats are filtered out at an early stage. An important building block of this layered security is mail server protection. By scanning all incoming and outgoing e-mail, malware and spam can be prevented from ever reaching their destination.

How to deploy mail server security depends on the current mail server setup. Even for companies that are not locally hosting a mail server, scanning all e-mail is recommended. This can be achieved by deploying MailSecurity as a gateway, collecting e-mail messages from an external POP3 server and scanning them before relaying them to local clients. Local mail servers can be covered by installing MailSecurity as a gateway on the mail server itself or on a dedicated mail security gateway server. When using Microsoft Exchange 2007 SP1, 2010, or 2013, deploy the MailSecurity Exchange plugin. Chapter 4.2.4 offers detailed information about the different deployment types and the installation procedure.

16.1. Exchange plugin

The Exchange plugin of MailSecurity complements existing Exchange workflows. By deeply integrating with the server, it provides transparent malware protection: without noticeable delays or user interaction, all objects are scanned and only passed on if they are free of malware. Plugin deployment looks like a regular client-server deployment. The Exchange plugin is deployed to the Exchange server and reports to a ManagementServer. This can be an existing network ManagementServer or a ManagementServer that is installed along with the Exchange plugin. The Exchange plugin will show up in G Data Administrator's client management area. When selecting it, an additional set of tabs becomes available. The **TASKS**, **REPORTS** and **STATISTICS** tabs offer functionality that is similar to their normal client management counterparts. On the **EXCHANGE SETTINGS** tab, status information can be checked and malware scan settings can be configured. The status information includes the engine and plugin version numbers and the last access date. In the menu bar, the **REFRESH** button can be used to refresh status information. In addition, it offers buttons to enable automatic virus signature and engine program file updates. When enabled, the Exchange plugin will automatically update itself whenever it connects to the ManagementServer (in the interval defined under **SERVER SETTINGS > SYNCHRONIZATION**). A manual update can be initiated at any time.

16.1.1. On-access scan

Several types of scans can be configured by MailSecurity for Exchange. The most important one is the on-access scan. Comparable to the file system monitor of G Data Security Client, the on-access scan monitors all actions on the Exchange server. If a client tries to access an e-mail, attachment or other object, it is scanned automatically and only made available if it is free of malware. Under **SCAN SETTINGS**, the **EXCHANGE SETTINGS** tab lets you configure scan parameters for the on-access scan. The scan can be carried out using one or two scan engines. Using two engines provides optimal security and is the recommended option. However, if scan performance is not as good as expected, one of the two engines can be disabled. This still offers very good detection while increasing performance. Scan performance can be further influenced by selecting the type of files that should be scanned. The most secure option is to scan all files, but this does take more time than a limited scan. A limited scan only includes program files and documents, the file types most likely to be infected. Heuristics

can be used to further increase detection by analyzing typical characteristics of malware. It slightly increases the chance of getting false positives, but greatly enhances malware detection. Enabling scans of archive files makes sure that even malware hiding inside of archives is found. This does increase scan time and if an infected file is found within the archive, the complete archive will be quarantined or removed. This option can be disabled if all clients are using the file system monitor, to ensure that malware is picked up as soon as it is extracted from the archive.

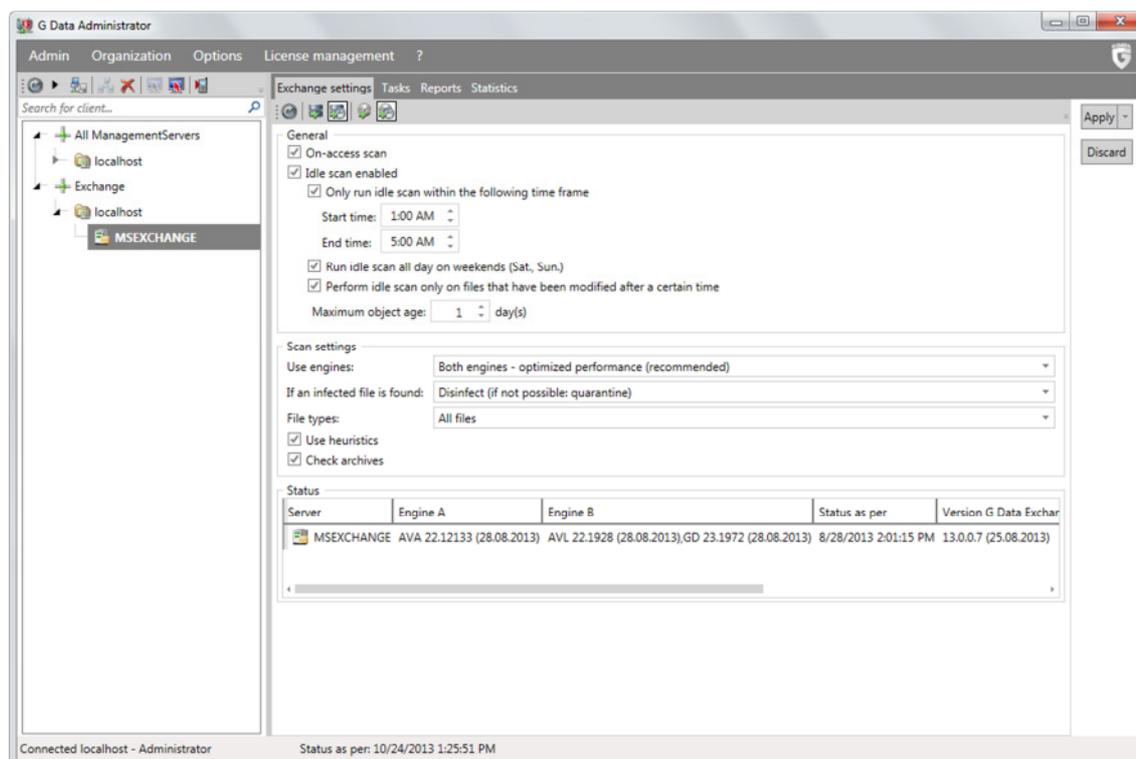


Image 63: G Data Administrator, Exchange settings

If malware is detected, several actions can be taken. The recommended option is to try to remove the malware from the file, moving it to the quarantine if removal does not succeed. This will prevent data loss as much as possible, while making sure the malware cannot be run. The REPORTS module will show a report when malware is blocked and allow you to examine quarantined files. Alternatively, MailSecurity can delete infected attachments, delete the entire message or log the threat without blocking it. Immediately deleting an infected object is the most secure option, but can cause data to be removed in case of a false positive detection. Only logging threats is not recommended; this will ignore any detected malware, allowing Exchange and its clients to access and possibly execute it. Even though a report will be added to the REPORTS module, letting administrators to manually take action, the time window between report and action allows for infection and potentially further distribution.

16.1.2. On-demand scan

In addition to the on-access scan, the TASKS module lets you schedule single and periodic scans, which function similar to client scan jobs. The settings are identical to those for client scan jobs, except for options that are not relevant for Exchange objects. Instead of defining the scan scope using the file system, an Exchange scan job is defined to operate on specific mailboxes. As with file system scans, it is recommended to scan all objects on the server regularly. This can be achieved by

planning a weekly, biweekly or monthly scan job. A full scan job can be very performance-intensive. It should be scheduled during off-peak times, such as the weekend or at night. Alternatively, the `IDLE SCAN` option can be enabled on the `EXCHANGE SETTINGS` tab. Like its client counterpart, the Exchange Idle scan is only run when Exchange Server is not doing anything else, preventing performance decrease. However, this does not include other processes on the server (for example, if it is also being used as a web server). To make sure that the Idle scan will never conflict with other server processes, it can be limited to run only within a specific time frame. For example, define a `START TIME` of 11:00 PM and an `END TIME` of 5:00 AM to run the idle scan only during the night. An exception can be made by selecting `RUN IDLE SCAN ALL DAY ON WEEKENDS`. To limit the Idle scan to files that have been modified recently (avoiding older, unused files), enter a `MAXIMUM OBJECT AGE`. This will reduce the amount of scan time but should be used with care: older files may remain infected if they are not scanned. The on-access scan should be enabled, to make sure that even infected older files are blocked.

16.2. MailGateway

MailSecurity MailGateway provides malware protection and spam filtering measures for all mail servers. It is compatible with all mail servers. When using Microsoft Exchange, malware protection is most easily implemented by installing the Exchange plugin of MailSecurity. In order to combine malware protection with spam filtering, deploy MailGateway instead of the plugin. For all types of mail servers, MailGateway can be integrated into the e-mail workflow by scanning traffic before it reaches the actual mail server. Various implementations are possible, such as an installation on the existing mail server or deployment to a dedicated mail gateway server. Chapter 4.2.4.2 details the possibilities. Just installing the gateway, however, is not enough. Mail traffic has to be routed through MailSecurity in order to be able to scan it. After completing this configuration, the individual mail traffic protection components (malware protection and spam filter) can be set up.

16.2.1. Deployment

After the installation wizard of MailSecurity MailGateway has been run, the MailGateway server is automatically started. Like ManagementServer, it carries out its tasks in the background without demanding user interaction. However, it requires some initial configuration in order to start scanning mail traffic for malware and spam. All settings for MailGateway can be edited using MailSecurity Administrator. The MailGateway setup wizard automatically installs MailSecurity Administrator on the same PC. However, like the Administrator application for ManagementServer, it can be used from any PC that can make a connection to the server on which MailGateway is running.

Launching MailSecurity Administrator for the first time, it will ask for the server name and password. The server name field will be filled automatically with the name of the server on which MailGateway has been installed. The password field can be left empty – clicking OK will prompt you for a new password. If the password needs to be changed at a later stage, click the button `Change password` on the `Advanced` tab of the `Options` window of MailSecurity Administrator. If the password has been lost and logging in to MailSecurity Administrator is no longer possible, the password can be reset by removing the key `HKEY_LOCAL_MACHINE\SOFTWARE\Wow6432Node\G DATA\AVKSmtplib\pw` (Windows 64-bits) or `HKEY_LOCAL_MACHINE\SOFTWARE\G DATA\AVKSmtplib\pw` (32-bits) from the Registry on the MailGateway server.

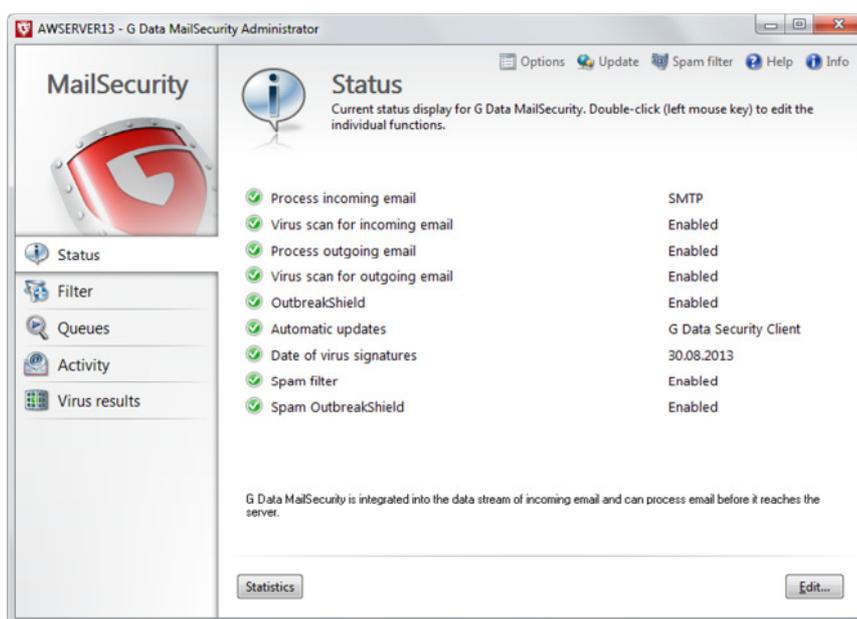


Image 64: G Data MailSecurity Administrator, Status

MailSecurity Administrator opens on the STATUS page, displaying information about the main functions of the gateway server. Each status message can be double-clicked to open the relevant part of the OPTIONS menu. When the program is started for the first time, some functions will show as being disabled. As part of the post-deployment tasks, automatic virus signature updates and mail server port settings will be configured. Subsequently, MailGateway's malware protection, spam filter, and custom filters can be set up.

16.2.1.1. Virus signature updates

An important first step is to configure automatic updates for the virus signature database. Click the UPDATE option in the top menu bar to open the INTERNET UPDATE window and choose one of two update methods. It is recommended to deploy G Data Security Client before deploying MailSecurity MailGateway (see chapter 4.2.4.2). Not only does this ensure that malware cannot infect the gateway server, it also allows MailGateway to scan incoming and outgoing e-mail with the Security Client's virus signatures, eliminating the need to maintain its own virus signature database.

If MailGateway is not running in a ManagementServer network, the alternative is to configure it to download its own virus signatures. To obtain updates, MailSecurity will log in to the online G Data UpdateServers, for which it requires access data. Click SETTINGS AND SCHEDULING to open the update settings and choose the tab ACCESS DATA. If user name and password have been obtained previously (for example, if the software has already been activated), they can be entered directly. Alternatively, click REGISTER WITH SERVER to enter your registration number. After registering with G Data, user name and password will be generated automatically. Make sure to write them down somewhere in order to be able to use them in case of a reinstallation. After entering access data, an update schedule should be defined on the INTERNET UPDATE SCHEDULE tab. To make sure that MailGateway can scan e-mails with the latest virus signatures, an hourly update should be defined. If the computer on which MailGateway has been installed uses a proxy server to connect to the internet, it should be defined on the INTERNET SETTINGS tab. This is also the place to enter login data, if required to set up an internet connection.

16.2.1.2. E-mail streams

It is important to distinguish between three possible e-mail streams to be scanned by MailGateway: incoming SMTP e-mail (from the internet into the network), outgoing SMTP e-mail (from the network to the internet), and POP3 (only for incoming e-mail). Not every network uses every type of e-mail communication. For the protocols that are in use, port configuration on both MailGateway and mail server is essential to route e-mail streams through MailGateway before delivering them to the mail server. Which ports need to be changed is determined by the deployment type.

Depending on the mail server, protocol and port settings may be hard to access. In various TechPapers, G Data offers additional configuration support for specific products. See the following TechPapers for more information: #0149 (Tobit David), #0150 (AVM Ken!), #0151 (Microsoft Exchange Server 2010), and #0152 (Microsoft Exchange Server 2007).

Deployment 1: On the mail server (with mail server port changes)

The advantage of deploying MailGateway to the mail server is that no separate server has to be configured; the downside is the need to carefully configure ports for SMTP and POP3. To avoid changes to the local firewall or DNS records, some mail server ports should be changed. An example configuration is the following: MailSecurity receives retrieval requests for POP3 on port 7110 (either from clients or from a POP3 connector on the mail server), connects to the internet POP3 mail server, collects e-mail, processes it and delivers it. Inbound SMTP mail (from the internet) will be received by MailSecurity on port 25, processed, and forwarded to the internal mail server on port 7125. Outbound SMTP mail will be received by the internal mail server on port 7125, forwarded to MailSecurity on port 7025, processed, and sent. The internal mail server will operate at SMTP port 7125. Exchange clients do not need to be reconfigured, but clients that send mail using SMTP settings have to be configured to connect to that port. The port settings for MailSecurity and the mail server looks as follows:

Traffic	MailGateway	Mail server
SMTP (incoming)	25	7125
SMTP (outgoing)	7025	7125
POP3	7110	110

To allow MailGateway to scan incoming SMTP e-mail, open the **OPTIONS** dialog and select the **INCOMING (SMTP)** tab. Under **RECEIVED**, enable the option **PROCESS INCOMING MAIL**. MailGateway operates at port 25 and the mail server uses a different port (for example, 7125). Once incoming SMTP e-mail has been scanned, it is delivered to the mail server according to the settings listed under **Forwarding**. By default, **USE DNS TO SEND EMAIL** is checked. This option uses the local **MX** record to determine where to deliver e-mail, which will only work if the **MX** record lists the mail server. Since MailGateway will be listed in the **MX** record in many cases, it is safer to disable **USE DNS TO SEND EMAIL**. The mail server IP and port can then be specifically defined. Enter *127.0.0.1* as IP address and enter the mail server port in the **Port** field (in this example *7125*).

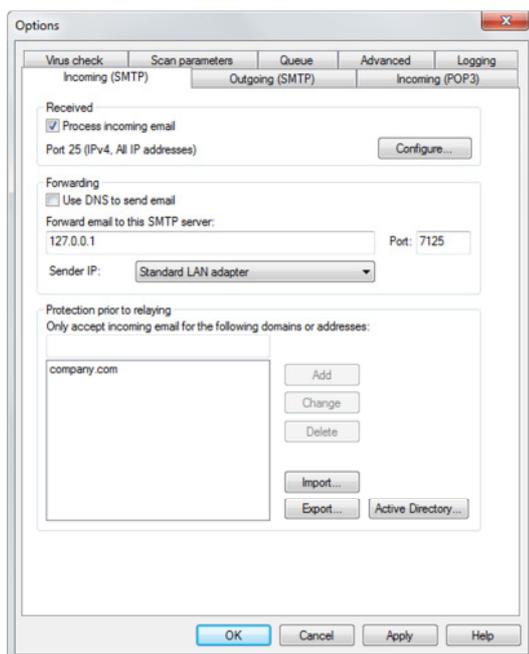


Image 65: G Data MailSecurity Administrator, Options, Incoming (SMTP)

To scan outgoing e-mail, open the OUTGOING (SMTP) tab. Enable PROCESS OUTGOING EMAIL and enter a port number. The simplest configuration is to define a port that is different from the INCOMING (SMTP) port, in this example 7025. However, if you decide to use the same port for incoming and outgoing SMTP traffic, MailGateway needs to be able to tell apart incoming and outgoing SMTP traffic. This can be achieved by adding 127.0.0.1 and the server's IP address (for example, 192.168.1.2) to IP ADDRESSES/SUBNETS OF COMPUTERS THAT SEND OUTGOING EMAIL. Outgoing e-mail can be delivered in two ways. MailGateway can directly deliver e-mail using the target's DNS record. This is the recommended configuration. If USE DNS TO SEND EMAIL is disabled, you can define an SMTP server to which outgoing e-mail should be forwarded. If any authentication is required, select the appropriate procedure in the AUTHENTICATION window.

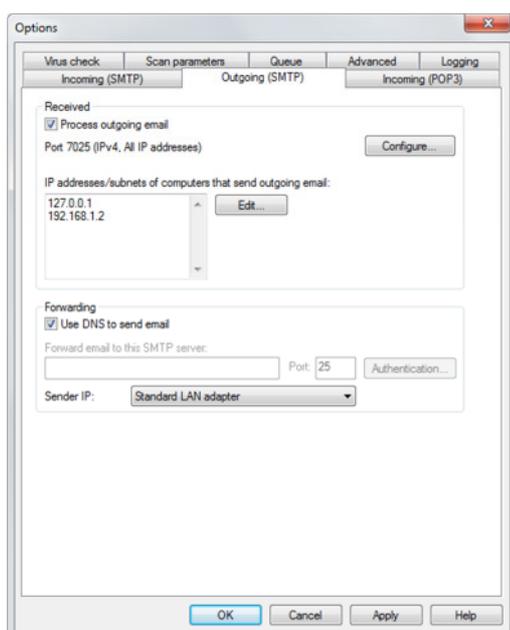


Image 66: G Data MailSecurity Administrator, Options, Outgoing (SMTP)

POP3 scans can be configured on the INCOMING (POP3) tab. Enable PROCESS POP3 ENQUIRIES and enter the port on which MailGateway will receive POP3 enquiries (7110). Tick PREVENT EMAIL PROGRAM TIMEOUT to prevent the recipient from getting a timeout error from their e-mail software if POP3 retrieval is taking too long. Under COLLECTION, enter the POP3 server from which e-mail should be collected. This is usually the POP3 server of the internet service provider. MailGateway will use the same login which the client is using in its POP3 enquiry. Under FILTER, you can define a replacement text for e-mail messages that are rejected by the malware protection or the spam filter.

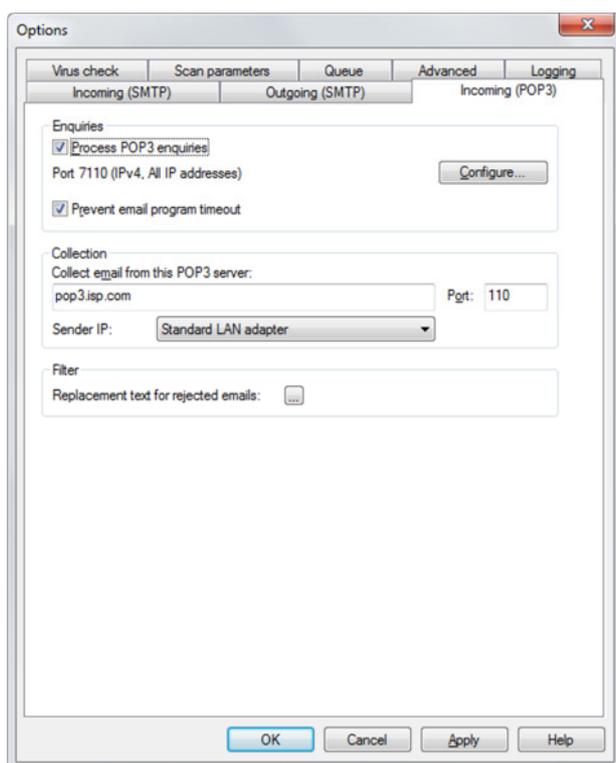


Image 67: G Data MailSecurity Administrator, Options, Incoming (POP3)

After having set up MailGateway, the mail server settings should be updated. Change the server for incoming POP3 e-mail to the server's IP address (192.168.1.2) and add the MailGateway port (7110). The port for incoming and outgoing SMTP e-mail should be changed to 7125. To properly route outgoing e-mail, change the mail server's setting for outgoing SMTP e-mail to the server's IP address and add the MailGateway port (7025). Disable any SMTP authentication, because it will be carried out by MailGateway.

Because the SMTP port of the mail server has been changed, clients that send e-mail using SMTP settings should send outgoing e-mail to the new port. Change the settings of local e-mail clients to reflect this (SMTP port 7125). Clients that send e-mail using Exchange do not need to have their settings updated.

Deployment 2: On the mail server (without mail server port changes)

If changes to the mail server ports are not feasible, for example if there are simply too many clients to reconfigure, the mail server can remain operational on port 25, and MailGateway can receive and send SMTP mail using ports 7025 and 7125 respectively. However, to ensure that MailSecurity will still be the first security layer, the local firewall needs to be updated to forward incoming SMTP traffic to MailGateway's non-standard port 7025. The port settings for this scenario look as follows:

Traffic	MailGateway	Mail server
SMTP (incoming)	7025	25
SMTP (outgoing)	7125	25
POP3	7110	110

To allow MailGateway to scan incoming SMTP e-mail, open the Options dialog and select the Incoming (SMTP) tab. Under Received, enable the option Process incoming mail. MailGateway operates at non-standard port 7025. Once incoming SMTP e-mail has been scanned, it is delivered to the mail server according to the settings listed under Forwarding. By default, Use DNS to send email is checked. This option uses the local MX record to determine where to deliver e-mail, which will only work if the MX record lists the mail server. Since MailGateway will be listed in the MX record in many cases, it is safer to disable Use DNS to send email. The mail server IP and port can then be specifically defined. Enter *127.0.0.1* as IP address and enter the mail server port in the Port field (in this example 25).

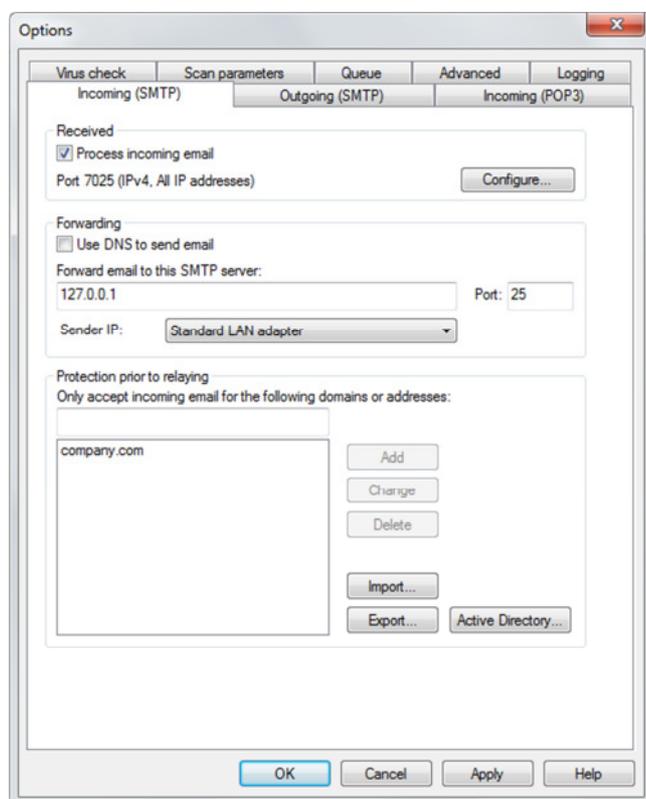


Image 68: G Data MailSecurity Administrator, Options, Incoming (SMTP)

To scan outgoing e-mail, open the OUTGOING (SMTP) tab. Enable PROCESS OUTGOING EMAIL and enter a port number. The simplest configuration is to define a port that is different from the INCOMING (SMTP) port; in this example 7125 is used. However, if you decide to use the same port for incoming and outgoing SMTP traffic, MailGateway needs to be able to tell apart incoming and outgoing SMTP traffic. This can be achieved by adding *127.0.0.1* and the server's IP address (for example, *192.168.1.2*) to IP addresses/subnets of computers that send outgoing email. Outgoing e-mail can be delivered in two ways. MailGateway can directly deliver e-mail using the target's DNS record. This is the recommended configuration. If USE DNS TO SEND EMAIL IS DISABLED, you can define an SMTP server to which outgoing e-mail should be forwarded. If any authentication is required, select the appropriate procedure in the AUTHENTICATION window.

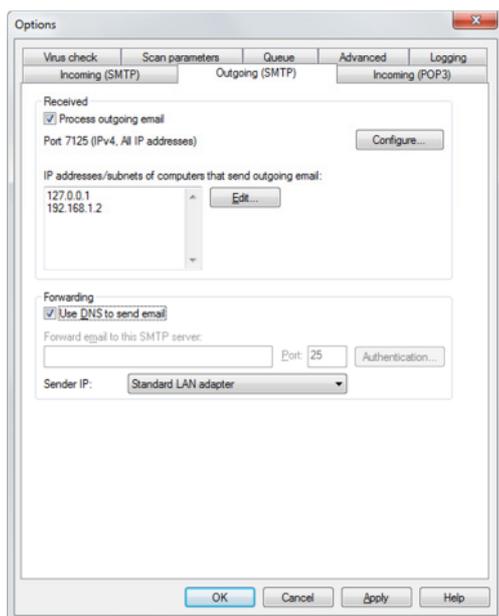


Image 69: G Data MailSecurity Administrator, Options, Outgoing (SMTP)

The settings for POP3 e-mail retrieval are identical to those of deployment 1.

The port settings for the mail server itself do not need to be changed, but it does need to be able to locate MailGateway. Change the server for incoming POP3 e-mail to the server’s IP address (192.168.1.2) and add the MailGateway port (7110). To properly route outgoing e-mail, change the mail server settings for outgoing SMTP e-mail to the server’s IP address and add the MailGateway port (7125). Disable any SMTP authentication, because it will be carried out by MailGateway. Because the mail server is still using port 25 for incoming and outgoing SMTP e-mail, a network-level change is necessary to make sure MailSecurity receives incoming SMTP e-mail. The router or firewall should forward incoming SMTP traffic on port 25 to MailGateway on port 7025.

Deployment 3: Dedicated gateway server

The alternative to an installation on the mail server is deploying MailGateway to its own server. Mail server ports do not need to be changed, but incoming e-mail needs to be delivered to MailGateway first, before it arrives at the mail server. This can be achieved in a couple of ways. The MX record in the DNS entry for the network’s domain can be changed from the mail server’s IP to MailGateway’s IP, redirection can be set up in the firewall, or the original IP address of the mail server can be assigned to MailGateway. In this example, the ports are configured as follows:

Traffic	MailGateway	Mail server
SMTP (incoming)	25	25
SMTP (outgoing)	7025	25
POP3	110	110

To allow MailGateway to scan incoming SMTP e-mail, open the OPTIONS dialog and select the INCOMING (SMTP) tab. Under RECEIVED, enable the option PROCESS INCOMING MAIL. MailGateway operates at port 25. Once incoming SMTP e-mail has been scanned, it is delivered to the mail server according to the settings listed under FORWARDING. By default, USE DNS TO SEND EMAIL is checked. This option uses the local MX record to determine where to deliver e-mail, which will only work if the MX record lists the mail server. Since MailGateway will be listed in the MX record in many cases, it is safer to disable USE DNS

TO SEND EMAIL. The mail server IP and port can then be specifically defined. Enter the mail server's IP address (in this example *192.168.1.3*) and enter the mail server port in the PORT field (in this example 25).

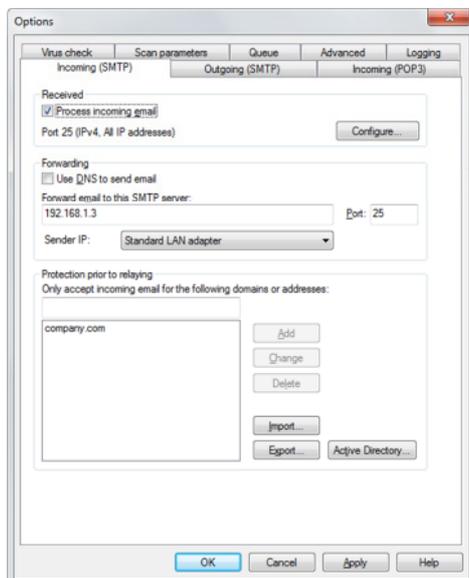


Image 70: G Data MailSecurity Administrator, Options, Incoming (SMTP)

To scan outgoing e-mail, open the OUTGOING (SMTP) tab. Enable PROCESS OUTGOING EMAIL and enter a port number. The simplest configuration is to define a port that is different from the INCOMING (SMTP) port, in this example 7025 is used. However, if you decide to use the same port for incoming and outgoing SMTP traffic, MailGateway needs to be able to tell apart incoming and outgoing SMTP traffic. This can be achieved by adding the mail server's IP address (for example, *192.168.1.3*) to IP ADDRESSES/SUBNETS OF COMPUTERS THAT SEND OUTGOING EMAIL. Outgoing e-mail can be delivered in two ways. MailGateway can directly deliver e-mail using the target's DNS record. This is the recommended configuration. If USE DNS TO SEND EMAIL is disabled, you can define an SMTP server to which outgoing e-mail should be forwarded. If any authentication is required, select the appropriate procedure in the AUTHENTICATION window.

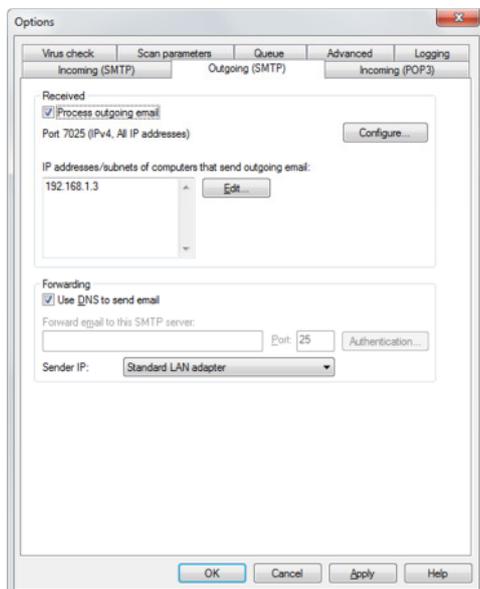


Image 71: G Data MailSecurity Administrator, Options, Outgoing (SMTP)

POP3 scans can be configured on the INCOMING (POP3) tab. Enable PROCESS POP3 ENQUIRIES and enter the port on which MailGateway will receive POP3 enquiries (110). Tick PREVENT EMAIL PROGRAM TIMEOUT to prevent the recipient from getting a timeout error from their e-mail software if POP3 retrieval is taking too long. Under COLLECTION, enter the POP3 server from which e-mail should be collected. This is usually the POP3 server of the internet service provider. MailGateway will use the same login which the client is using in its POP3 enquiry. Under FILTER, you can define a replacement text for e-mail messages that are rejected by the malware protection or the spam filter.

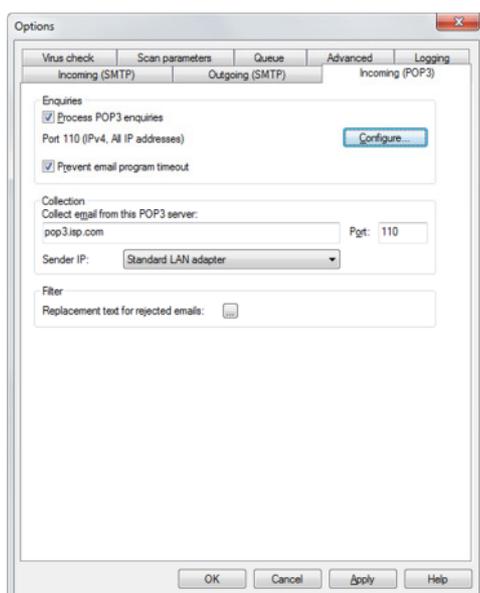


Image 72: G Data MailSecurity Administrator, Options, Incoming (POP3)

After setting up MailGateway, the mail server settings should also be updated. Change the server for incoming POP3 e-mail to the MailGateway's IP address (192.168.1.2) and add the MailGateway port (110). To properly route outgoing e-mail, change the mail server settings for outgoing SMTP e-mail to the MailGateway's IP address and add the MailGateway port (25). Disable any SMTP authentication, because it will be carried out by MailGateway.

The MailGateway server should replace the mail server to receive incoming SMTP e-mail. This can be achieved by configuring the router or firewall to redirect incoming SMTP traffic to the MailGateway server (192.168.1.2). Alternatively, the DNS MX record can be changed to MailGateway's (external) IP address. Finally, the IP address of the mail server can be assigned to MailGateway, and a new IP address can be assigned to the mail server.

16.2.1.3. SMTP relay protection

When enabling the processing of incoming SMTP e-mail, MailGateway should be protected against relay abuse. On the INCOMING (SMTP) tab, under PROTECTION PRIOR TO RELAYING, define domains for which MailGateway accepts e-mail. If no domain is added, no incoming SMTP e-mail is accepted! By adding only the company domain(s), e-mail for other domains is automatically dropped, making sure spammers cannot use the MailGateway server to distribute unwanted e-mail. If incoming e-mail should be accepted for all domains, add the domain *.*. As an alternative to adding domains manually, Active Directory can be used to import a list of addresses. E-mail to addresses that are not on the list is automatically dropped. Click ACTIVE DIRECTORY to configure AD server settings and update interval.

16.2.1.4. Multiple POP3 servers

If e-mail should be collected from more than one POP3 server, the standard deployment scenario needs to be slightly altered. Instead of defining one POP3 server in the MailGateway settings, the POP3 server can be defined for each account on mail server (when collecting POP3 using Exchange, for example) or client (when collecting POP3 e-mail on clients). In MailSecurity Administrator, open the INCOMING (POP3) tab of the OPTIONS window. Leave all existing settings as-is, but remove the POP3 server address that has been entered under COLLECTION. In the mail server or local client settings, define an account for each POP3 server that needs to be checked. As POP3 server, enter the IP address and port of MailGateway. Edit the user name field for each account, adding the POP3 server name in front of the user name, separating both values by a colon. For example, to check the account for user name *company* on server *pop3.isp.com*, enter *pop3.isp.com:company*. The password field should contain the account password.

16.2.1.5. Queue

MailGateway processes e-mail messages immediately upon receipt. Malware scans and spam filtering are carried out and the message is forwarded to the mail server according to the deployment settings. The receiving mail server, however, may not always be reachable. E-mail messages that cannot be delivered end up in the queue. MailGateway will repeatedly retry to deliver queued messages for a certain amount of time. To view the messages that are currently queued, use the QUEUES module of MailSecurity Administrator. Using the INCOMING/OUTGOING button, you can switch between viewing the queue of incoming and outgoing messages. For each message, the target host, sender, next repeat, and status are shown. No action is required for any of the queued messages; MailGateway automatically takes care of them. Optionally, MailGateway can try to deliver the message immediately by clicking REPEAT NOW. Using the DELETE button, queued messages can be permanently deleted from the queue, preventing delivery.

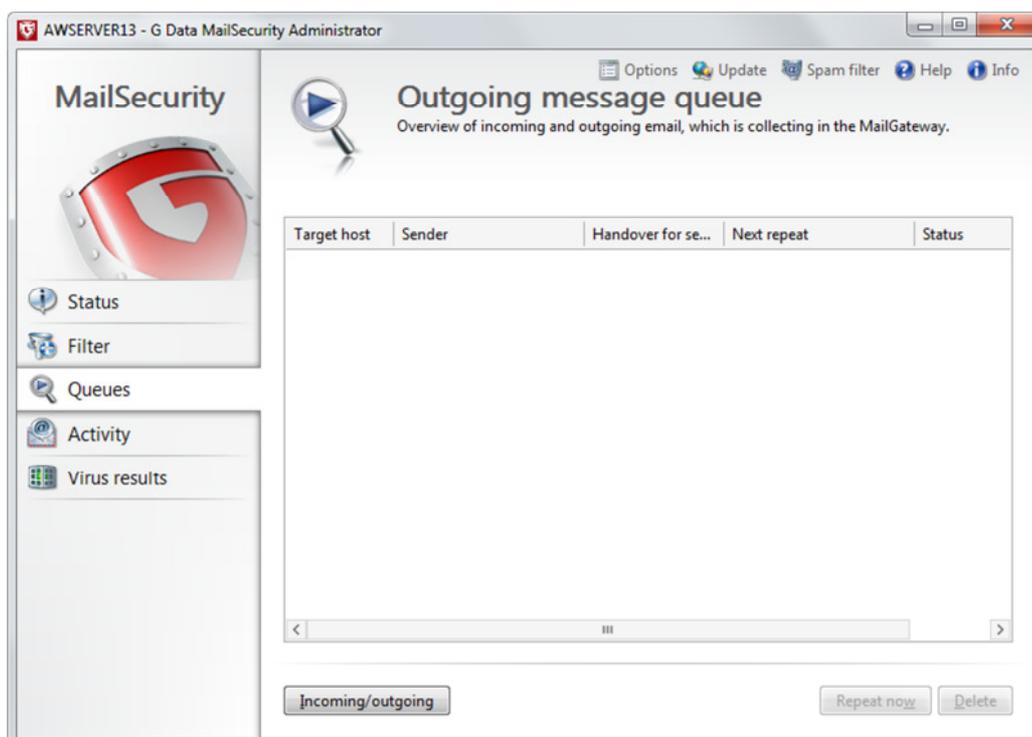


Image 73: G Data MailSecurity Administrator, Queue

MailGateway's handling of queued messages can be configured in the Options window on the Queue tab. It will regularly try to deliver messages that are in the queue up to a maximum amount of time. Under Error waiting time, this maximum can be defined. The default value is 24, which means that the message will be discarded if it still has not been delivered after 24 hours. Under Repeat interval, retry intervals can be defined. MailGateway will try to deliver queued messages in these intervals until the maximum amount of time is reached. The default values are 0.25, 0.5, 1, 4. This means that MailGateway's first retry will be after 15 minutes, the second will be half an hour later, the third will be an hour later, and then every 4 hours until the maximum. The default values can be changed to fit individual needs. For example, retries can be scheduled with shorter intervals (for example, 0.25 to retry every 15 minutes until the maximum is reached). This will prevent further delays if a mail server was only offline temporarily, but can cause an increased amount of mail traffic. In any case, e-mail senders should be informed if their message cannot be delivered (yet). They can be notified every few hours (the default value is 4). Avoid sending too many notifications to senders – only if the maximum waiting time is changed, the notification interval should be changed proportionally. Finally, the number of messages in the queue can be limited. This option is disabled by default, because limiting the queue leads to message loss if the maximum is reached: e-mails that cannot be delivered are deleted immediately if there are too many messages in the queue already. Only if the number of messages in the queue inhibits performance or otherwise causes problems should the queue size limit be enabled.

16.2.2. Administration and migration

Having deployed and configured malware protection, spam filter and custom filters, MailGateway carries out its scans and filtering without intervention. No confirmation is needed for MailGateway to reject or remove e-mail messages. However, it is recommended to monitor MailGateway's performance and actions, to make sure that no unexpected actions are taken. The most direct method to monitor MailGateway's handling of spam is to configure the forwarding of spam messages (see chapter 16.2.4). On a more general level, administrators can use MailSecurity Administrator to review MailGateway's actions. The ACTIVITY module shows a log of actions carried out by MailGateway, such as e-mail messages being received, processed or sent. If the database for statistical assessment has been installed (see chapter 4.2.4.2), the STATUS module shows a STATISTICS button. The statistics window shows general statistics about processed e-mail messages, as well as top 10 lists for spam addresses, spam IP addresses and detected viruses. The statistics can be shown for each of the traffic streams that MailGateway analyzes (SMTP incoming, SMTP outgoing and/or POP3 incoming). To configure the statistics, open the LOGGING tab of the OPTIONS window. Select SAVE IN THE DATABASE to save statistics about e-mail traffic in the database, enabling analysis. Alternatively, some details about the traffic can be saved to a log file (saved in MailGateway's installation folder, by default C:\Program Files (x86)\G DATA\G DATA MailSecurity\maillog.txt). The log file contains timestamp, sender and recipient, mail size and spam value (see chapter 16.2.4.1). To limit file size, logging can then be limited to junk mail only, or to a maximum number of e-mails.

In case MailGateway is being migrated to another server, its settings can easily be exported. In the OPTIONS menu, the ADVANCED tab offers the button EXPORT. It allows administrators to save an XML file containing MailGateway's settings. On the new server, it can be imported using the IMPORT button. Similarly, existing custom filters (see chapter 16.2.5) should be exported. The IMPORT and EXPORT buttons in the FILTER module can be used to manage filter settings as XML files. If the Bayesian content filter has been enabled (see chapter 16.2.4.5), its configuration needs to be imported separately to retain its progress. Using the Windows SERVICES panel, stop the MailGateway service on

the old and new MailGateway server. The configuration file is located in MailGateway's FILTER folder (by default: C:\Program Files (x86)\G DATA\G DATA MailSecurity\filter). Copy the files BayesSpam.txt and BayesNotSpam.txt into the FILTER folder on the target server and restart the MailGateway service.

16.2.3. Malware protection

E-mail remains a very popular attack vector for criminals. Many forms of malware are distributed via e-mail, as part of generic spam runs as well as targeted attacks trying to gain access to a company network. To make sure that malicious programs do not end up infecting network clients, incoming e-mail should be cleaned of malware. At the same time, e-mail that is being sent by network clients should also be checked. Infected clients may unknowingly send malware-infected e-mails, putting external e-mail recipients in danger. MailGateway scans incoming and outgoing traffic for malware using the same scan technology that is included in its Security Client. Using two scan engines, malware is detected before it reaches the recipient. Various measures can be carried out, such as disinfecting, renaming or deleting the infected object.

16.2.3.1. Inbound and outbound e-mail

The OPTIONS window of MailSecurity Administrator can be used to configure malware detection. On the tab VIRUS CHECK, parameters for scanning inbound and outbound e-mail can be defined. It is recommended that both inbound and outbound e-mail streams are checked for malware. For inbound e-mail, the most important choice is whether to deliver infected messages or not. The recommended option is to try to disinfect infected messages or attachments, and deleting them if disinfection is not possible. This will prevent unnecessary data loss while ensuring that infected content will never be delivered to end users. Alternatively, infected attachments or messages can be directly deleted. This is the most secure measure, but may lead to data loss if an object is incorrectly identified as malware. Renaming infected attachments is the third option, but is not recommended. While this does prevent direct access to malware, end users could still get infected if the files is renamed back and opened. To notify the recipient of the e-mail message that malware has been found, a virus alert can be added to the subject and body of the infected message. This is recommended to make sure that the recipient is aware of the fact that the message has been altered and is no longer the message as it was originally sent – emphasizing the fact that the virus has been removed and infection is no longer possible. If the e-mail message includes a password-protected attachment, the attachment cannot be scanned. To warn the recipient, a notice can be inserted into the message body, explaining that one or more attachments have not been scanned. In addition to the recipient notice, you can define one or more e-mail addresses (separated by semicolons) that should receive an alert when malware is found. For standalone MailGateway installations, this is an easily configurable way to make sure the administrator is informed. If MailGateway has been deployed in the same network as an existing ManagementServer installation, it is recommended to use the option REPORT VIRUS RESULTS TO G DATA MANAGEMENTSERVER instead.

For outbound mail, the security options are slightly different. Although outbound messages are scanned by the same scan engines as incoming messages, the reaction to infected messages is different: they are not sent. Disinfection or renaming attachments is not possible, as delivery of messages with possible malware should be prevented at all cost. If malware is detected, the sender can be notified. It is recommended to send a notification to explain that the message will not be delivered. If malware is not detected, MailGateway can attach a report to outbound messages, including version information. This will inform the recipient that the message has been scanned for

malware before it was sent. As with inbound messages, you can define e-mail addresses to which an alert should be sent in case of a malware infection.

16.2.3.2. Scan parameters

The settings for the malware scan can be customized on the `SCAN PARAMETERS` tab. The default configuration provides optimal protection. Only if problems with performance or with one of the individual protection modules arise should any of the scan parameters be changed. If the scan process with two engines is too demanding, one of the engines can be disabled. This will provide better performance but will slightly decrease the level of security. Scans can be limited to certain file types, such as program files or documents, or a custom list of extensions, but the safest option is to scan all files (enabled by default). Heuristics are also enabled by default, providing additional security by recognizing unknown malware based on typical characteristics. Archives can be excluded from scans by disabling `CHECK ARCHIVES`. If a file within an archive is found to be infected, the whole archive is renamed or removed. To prevent accidental data loss, checking archives for malware can be disabled, but this will lead to decreased security. OutbreakShield provides protection against mass mailings containing malware, even before virus signatures become available, and should be enabled in most cases. It does not disinfect e-mail messages that contain malware. Under `SETTINGS`, a replacement text for the mail body can be entered, so that recipients will know that an e-mail was blocked by OutbreakShield. Finally, phishing protection blocks e-mails that try to obtain passwords, credit card numbers, or other personally identifiable information by posing as an email from a legitimate institution.

16.2.4. Spam filter

A large percentage of e-mail traffic consists of spam. While not containing malware, many of these messages are unwanted, such as mass mailings of pharmaceutical ads or illegal software sales. On the client level, spam filters have long been deployed in order to remove incoming spam messages before they reach the inbox. This is an effective way to make sure that individual users do not need to spend time on reading and removing them, but it requires each client to have local spam filtering capabilities, such as a network- or client-specific configuration, an up-to-date rule set of spam definitions and local self-learning capabilities. MailGateway centralizes spam filtering capabilities, removing spam from e-mail traffic before it is even delivered to the clients.

MailGateway's fully automatic spam filter can be configured by clicking on the `SPAM FILTER` button in the top menu bar. On the `FILTER` tab, it can be enabled or disabled as a whole; by default, it is enabled. Add a name and note to identify the filter in the list view of the `FILTER` module (see chapter 16.2.5). The spam filter scans all e-mails messages and assigns them one of four categories: not spam, suspected spam, high spam probability, and very high spam probability. E-mails that are classified as not being spam are delivered immediately. For the other three categories, custom reactions can be defined. On the `FILTER` tab, click `CHANGE` to edit reaction settings for each category. E-mail messages can be rejected, a spam warning can be added to their subject or body, the message sender can be notified, and the mail can be forwarded to someone. It is safe to configure MailGateway to reject e-mails that have a high or very high spam probability. Messages can optionally be forwarded to an administrator, who can fine-tune the filter if messages are not categorized correctly. Suspected spam should not be rejected immediately – as with malware scans, there is a risk of false positives. Since only a small number of messages falls in this category, they can be delivered to the recipient. As a warning, a subject prefix or body message should be added to the message.

16.2.4.1. Value-based ranking

Having configured the measures to be taken when MailGateway detects spam, the parameters that govern that detection can be tweaked. At the core of the automatic spam filter is a value-based ranking system which categorizes a message as spam as soon as it reaches a certain value. Individual criteria add points to a message's spam index value, such as its subject, body or (absence of) message ID. Criteria like mail size or content filter score subtract points. After adding the different values together, the final score determines whether the message is considered spam or not. On the Advanced settings tab, the individual scoring components can be adjusted. This should hardly ever be necessary – only if spam is incorrectly marked safe or if messages are accidentally treated as spam. In addition to the ranking system, messages can be marked as spam by one of the various other spam filter parameters.

16.2.4.2. Blacklist/Whitelist

Using the BLACKLIST tab, e-mail addresses or domains can be defined as spam. Any incoming e-mail message from one of the listed addresses or domains is immediately marked as spam. The WHITELIST tab does the opposite: whitelisted e-mail addresses and domains can always send e-mail messages; they will never be marked as spam. Combining the blacklist and whitelist allows administrators to fine-tune the spam filter. If a specific type of message, such as a recurring newsletter, is consistently marked as spam, it can be whitelisted. Conversely, if a recurring e-mail message is unwanted, it can be explicitly blacklisted.

16.2.4.3. Keywords

Instead of categorizing an e-mail by its sender address or domain, MailGateway can scan it for specific keywords. If any of them appear in a message, it is directly marked as spam. Separate keyword lists can be defined for message subject and body. By default, both lists are enabled and pre-populated with terms that are commonly used by spammers. If specific content should always be filtered as spam, keywords can be added to either or both of the lists. By default, the option MATCH WHOLE WORDS ONLY is enabled. This makes sure that a message is not categorized as spam if it contains a word which partly matches a keyword. For example: to filter out messages involving cash, the keyword "cash" can be added to the subject and body keyword lists. With the option MATCH WHOLE WORDS ONLY disabled, MailGateway would detect all instances of "cash", even in words like "cashew" or "cashier". To avoid this kind of unintentional filtering, the option should remain enabled.

16.2.4.4. Real-time blacklists

Real-time blacklists (RBLs) can be configured on a dedicated tab. RBLs are typically managed by online anti-spam organizations and contain lists of mail servers known to be used by spammers. When MailGateway is configured to use real-time blacklists, it will consult a selection of online lists to see if the sender's domain has been blacklisted. If it has been blacklisted, points are added to the message's spam index value. MailGateway can use the RBLs that have been defined by default, but they can be replaced with other URLs if required. If a domain should always be allowed, regardless of whether it is on an RBL, it can be whitelisted.

16.2.4.5. Bayesian content filter

Having defined all anti-spam parameters, the `CONTENT FILTER` can be enabled to add an additional layer of protection. It is a self-learning filter which relies on the principle of Bayesian spam filtering. When enabled, characteristics of every e-mail that has been designated as spam are stored in a database. Based on these characteristics, future e-mails are scanned. This helps detect spam messages even if they are slightly different than what has been defined using parameters such as whitelist, blacklist and keywords. The `CONTENT FILTER` tab shows statistics about the amount of e-mail that has been added to the filter database – the more e-mail is scanned, the better the Bayesian filter functions.

16.2.5. Custom filters

Malware and spam pose threats to e-mail recipients, but they are not the only kind of content that warrants filtering. In many corporate environments, e-mail is only meant to be used for business-related communication. Inappropriate, illegal or distracting contents should be filtered out. In addition to its protection against malware and spam, MailGateway offers several custom filters that can be used to filter out e-mail messages that do not meet company policies or are otherwise redundant.

To manage custom filters, open the `FILTER` module of MailSecurity Administrator. It provides an overview of all current filters. Under the list, a group of buttons offers functionality that applies to all filters. Settings can be imported and exported for all filters at once (see chapter 16.2.2). Using the `STATISTICS` button, some basic statistics can be shown for each filter: the number of e-mail messages that has been processed and the number that matched the specific filter. To manage filters, use the `NEW`, `EDIT` and `DELETE` buttons. To enable or disable a filter, use the filter's checkbox.

Each filter has specific options, which can be edited upon adding the filter or by clicking the `EDIT` button. All filters have a `NAME` and `NOTE` field. The filter name can be edited to tell filters apart in the filter list. This is especially useful for filters that can have multiple instances, such as the `SENDER` and `RECIPIENT` filter. The `NOTE` will also be displayed in the list – a useful place for comments about the specific filter settings.

Similar to the measures that can be configured for malware-infected files, MailGateway can take several actions when an e-mail message matches a filter. Most filters can be configured to reject messages that match the filter. Additionally, the message sender can be notified with a customizable text which informs the sender that the recipient has not received the message. An alert can be sent to one or more persons – a powerful measure for administrators that want to keep track of the status e-mail traffic in the network. Alerts can include the e-mail message that matched the filter. While this may be useful for analysis, messages may include malware-infected or otherwise inappropriate content and any forwarded messages should be handled with care.

Most filters can be applied to incoming as well as outgoing e-mails (except in cases where the filter logically only applies to one traffic type, such as the greylist filter). To provide safety for end users within the company, enabling filters for incoming traffic is recommended. Enabling one or more filters for outgoing e-mail makes sure that outgoing corporate e-mail will meet safety requirements.

16.2.5.1. Attachments

Attachments have long been a primary way of spreading malware. Early malware would distribute itself by sending an e-mail to all contacts in the victim's address book with a malicious executable

attached. Contemporary malware authors often execute targeted attacks, e-mailing only a select group of recipients with malware in the hopes of gaining access to valuable documents on their PCs. While malware will usually be picked up by the target PC's file system monitor, the risk can be even further decreased by having MailGateway filter risky attachments.

The attachment filter can be run in blacklist or whitelist mode. In blacklist mode, attachments are not allowed if they are on the list. In whitelist mode, only the listed attachment types are allowed; all others are removed. As with most filtering modules, using a whitelist mode provides the largest amount of security. By disallowing everything except previously approved extensions, even unknown attack vectors are blocked. However, legitimate content may also be blocked if it is of a file type that has not been explicitly defined as safe. Blacklist mode can be used to explicitly prohibit all known dangerous file types. It protects against the most frequently used malware attachment types, but may let malware pass if it uses an attack vector that has not been explicitly defined as unsafe.

The attachment types to be filtered can be entered in FILE EXTENSIONS list. Use semicolons to separate multiple extensions. When using the filter in blacklist mode, extensions to block commonly include executables, such as .exe, .scr and .com. Scripts can have extensions such as .bat, .vbs, .js and .cmd and should also be blocked. Registry setting files (.reg) can also be blocked. Archive files, such as .zip, .rar and .7z, may contain malware and can be filtered. Any other extension may be blocked, for example if it does not meet company policies. When whitelisting extensions, the list should be based on the most commonly received file formats, such as .txt, .doc or .jpg. Whitelisting attachments does not necessarily mean that they are safe. Documents could still contain malware. MailGateway's malware protection should scan the attachments and clients should enable e-mail scan or file system monitor. By enabling the option ALSO FILTER ATTACHMENTS IN EMBEDDED EMAIL, MailGateway even filters attachments if they are nested within an embedded e-mail file. This option should remain enabled, to make sure that even messages that are being forwarded as an embedded mail file are secured.

When MailGateway detects one or more offending attachments in an e-mail message, it can take multiple measures. When ONLY RENAME ATTACHMENTS is enabled, the attachment is renamed by adding the value added in the SUFFIX field. This is useful to prevent executables (including Office documents) from being run, as users would have to save and rename the file in order to run it. For additional security, uncheck ONLY RENAME ATTACHMENTS – MailGateway will then delete any attachment that matches the extension list (blacklist) or any attachment that does not match the list (whitelist). Removing the attachment is recommended, but may lead to data loss. If all client PCs are protected by a file system monitor, attachments can be renamed instead of removed, but this is not recommended. In addition to renaming or removing the attachment, a message can be inserted in the mail body. By adding a message, recipients can be informed that a virus was found – clearing any confusion that may arise from attachments being referred to in the message body.

16.2.5.2. Content

The content filter can be used as an extension of MailGateway's built-in spam filter. While the latter has been optimized for spam detection, including dedicated word lists and spam-centered modules, the content filter can be customized to filter any type of content. For example, content that is not related to work can be filtered out by adding filters for popular off-topic expressions. Another example is filtering sensitive content that could benefit competitors.

The type of content to be filtered can be entered in the form of a regular expression. Regular expressions are very powerful tools to match complex character sequences. Using the NEW button, a

regular expression can be built by entering several search terms that will be matched. This provides help for administrators that are not familiar with regular expressions. Alternatively, online resources about regular expressions can provide guidance¹⁷. The SEARCH SCOPE can be defined to include e-mail HEADER, SUBJECT, MAIL TEXT, HTML TEXT and/or EMBEDDED EMAIL.

16.2.5.3. External references

E-mails often contain HTML references to external content, such as images, links or scripts. While external content can be useful, for example for e-mail layout templates, it is often used by spammers and malware distributors. Using external images, spammers can gauge e-mail account activity: by embedding individualized image links into e-mails, a spammer can observe the e-mail client downloading the image, thus knowing that the e-mail has been read. In addition, the client IP address could be logged. To prevent external content from being downloaded at all, external references can be completely filtered out of incoming e-mails. The filter has no further options.

16.2.5.4. Greylist

The greylist filter takes advantage of the fact that spammers usually do not make use of an e-mail queuing system: a spam message is usually only sent once – the spammer’s mail server does not try to resend the message. Legitimate mail servers try to resend e-mail messages if they cannot be delivered on the first attempt. When the greylist filter is enabled, MailGateway does not immediately accept incoming e-mails. Instead, it shows the sending server a request to resend the message. Legitimate mail servers will honor the request and resend the message – spam servers will not. If the message is resent, it passes the greylist filter. Additionally, the combination of sender address, recipient address and sending mail server identification is then greylisted, to make sure that future e-mail messages are delivered immediately.

In order to enable the greylist filter, the spam filter must be active (see chapter 16.2.4). In addition, Microsoft SQL Server 2008 SP3 Express must have been installed during the installation of MailGateway (see chapter 4.2.4.2).

Upon enabling the greylist filter, several options can be defined. Under WAITING TIME, the amount of time for which an e-mail is held back can be edited. The default is 0 minutes for normal e-mails (the greylist is effectively turned off) and 30 minutes for e-mails that have been classified as suspicious. A message is deemed suspicious if the e-mail header looks like it has been tampered with, for example if a reverse lookup reveals discrepancies between mail server name and IP address. A suspicious message will only be delivered if the sending server resends the message after at least 30 minutes. If the sending server resends the message before the 30 minutes have expired, it is told to try again later.

To keep the greylist up to date, the combination of sender address, recipient address and sending mail server identification is only greylisted for a specific time. Two separate whitelist LIFETIMES can be set: with or without e-mail exchange. The value WITHOUT EXCHANGING EMAIL applies to greylisted e-mails that are not resent. If the sending mail server does not resend the e-mail within this amount of time (default 2 days), it is removed from the greylist. This measure makes sure that the greylist is not filled with entries from spam servers that send an e-mail only once without retrying. The value WITH EMAIL EXCHANGE governs greylist entries for which the sending mail server has resent the e-mail. Every

¹⁷ A starting point is the Wikipedia entry for „Regular expressions“, which includes a set of examples: http://en.wikipedia.org/wiki/Regular_expression#Examples.

time an e-mail is received that matches the greylist entry, it is delivered immediately and the greylist timer is reset. This ensures swift delivery for recurring e-mails, such as newsletters.

16.2.5.5. HTML scripts

Although HTML in e-mails can be used legitimately, for example to define layout, there is a security risk. Malicious scripts can be run upon opening an e-mail message, infecting a PC with malware or otherwise exhibiting unwanted behavior. The HTML script filter will remove all scripts from incoming and/or outgoing e-mail. Every HTML script tag pair (`<script>` and `</script>`) is filtered out, removing possible threats. The filter has no further options.

16.2.5.6. IP addresses

The IP filter can be used to blacklist or whitelist mail servers. When certain servers should not be allowed to send e-mail to the corporate network, they can be blacklisted by adding their IP address to the IP filter. Conversely, if only a few specific mail servers should be able to send e-mail, their IP address can be whitelisted. IP addresses can be entered as single addresses or using CIDN notation. The list of IP addresses can be exported as a simple text list and imported in case of a mail server migration or import from another system.

16.2.5.7. Language

If the network clients generally do not communicate using a specific language, it can be assumed that incoming messages in that language are unwanted. Using the language filter, messages that have been written in a specific language can be filtered. For example, an English-speaking company that does not have business partners or customers in Japan could block e-mails in the Japanese language to cut down on spam.

MailGateway's language recognition system assigns a language recognition rate to every e-mail. Under `UNDESIRED LANGUAGES`, the languages for which e-mails should be filtered out can be selected. In addition to the standard measures, MailGateway can add a spam warning to e-mail subject and body. The default subject prefix is `[%L %P]`, adding the language and language match percentage to the subject. Any message can be added to subject or body.

16.2.5.8. Read receipt requests

For many e-mail users, the practice of requesting read receipts provides certainty that the receiving party has actually read an e-mail message. For recipients, on the other hand, read requests can be an annoyance, especially if the e-mail warrants a reply anyway. Moreover, spammers may request read receipts to gauge activity on an e-mail account, increasing spam volume if the e-mail account is actively being used. The filter has no further options.

16.2.5.9. Sender/Recipient

E-mail messages can be filtered based on sender or recipient. If certain domains are sending unwanted e-mails that are not taken care of by any of the other filters, the domain can be added as a Sender filter to make sure it is filtered out. The same goes for e-mails that are addressed to specific recipient addresses or domains. The Recipient filter may not be used often in the fight against

unwanted e-mails, but can be set up to send an alert when e-mails are delivered to that specific recipient.

The Sender filter can be configured to match e-mails that have no sender, a practical measure to thwart spammers that manipulate an e-mail header to leave out sender information. The Recipient filter optionally matches messages with an empty recipient field (messages that only contain CC or BCC recipients).

16.2.5.10. Spam

The spam filter in the `FILTERS` module provides a shortcut to the extensive options window of the spam filter module (see chapter 16.2.4).

17. Advanced configuration

G Data business solutions have been designed to be ready for deployment out-of-the-box. No advanced configuration is required to get the server and client components up and running. Through the interface of G Data Administrator, all common settings are available. There are several advanced configuration possibilities available to make G Data solutions run optimally in an enterprise network. These settings do not usually need to be set and should only be changed if G Data Support asks you to do so.

17.1. Configuration files

Many advanced settings for G Data ManagementServer can be edited in XML-based configuration files. The two central files are Gdmms.exe.config and IUpdateCfg.xml, both located in the installation folder of G Data ManagementServer (typically C:\Program Files (x86)\G DATA\G DATA AntiVirus ManagementServer). The files can be edited using Notepad or any other text-based editor. When editing the configuration files, make sure that the text editor does not save the files with the .txt extension. Select ALL FILES (*.*) as file type and make sure the file name is exactly the same as it was before.

To avoid unwanted behavior, the background service of G Data ManagementServer should be shut down while editing configuration files, and restarted after saving them (use the SERVICES window to do this: Start, Run, *services.msc*). As with any advanced configuration, make sure that configuration files have been backed up before attempting any reconfiguration. If any of the components of the G Data solution exhibit unexpected behavior after the configuration change, revert to the backup. Do not edit any options other than the ones that apply to the problem that should be fixed.

```
<?xml version="1.0" ?>
<configuration>
  <configSections>
    <sectionGroup name="applicationSettings" type="System.Configuration ...
      <section name="Gdata.AntiVirus.Business.Server.Gdmms.Properties ...
    </sectionGroup>
  </configSections>
  <applicationSettings>
    <Gdata.AntiVirus.Business.Server.Gdmms.Properties.Settings>
      <setting name="UseIntegratedDatabase" serializeAs="String">
        <value>False</value>
      </setting>
      <setting name="DbDatabase" serializeAs="String">
        <value>GDATA_AntiVirus_ManagementServer</value>
      </setting>
    ...
  </applicationSettings>
</configuration>
```

Image 74: Gdmms.exe.config

The structure of Gdmms.exe.config will be familiar to administrators that have edited XML-style configuration files before. The file defines a section of application settings, and lists settings within the section with the tag pair <setting></setting>. Each setting has a name, defined as an attribute within the <setting> tag, and a value, enclosed in a <value></value> tag pair. If there is no value, the tag pair is shortened to <value />. Values are typically integer numbers, Boolean values (True/False), strings, or TimeSpan values. The tables in this chapter list the setting name and the corresponding possible values. The setting name should never be edited – only change the value to the desired setting.

The file IUpdateCfg.xml contains settings related to client and server program file as well as virus signature database updates. Its structure is slightly different: all settings are saved in their own tag pairs. To update settings in IUpdateCfg.xml, edit the value between the tag pair.

```
<?xml version="1.0" encoding="UTF-16" ?>
<IUpdate>
  <BasePath>C:\ProgramData\G DATA\AntiVirus ManagementServer\Updates</BasePath>
  <LoadDiffFiles>1</LoadDiffFiles>
  <UpdateCnt>5</UpdateCnt>
  <Language>GER</Language>
  <StandAlone>False</StandAlone>
</IUpdate>
```

Image 75: IUpdateCfg.xml

17.1.1. Database

During its installation, G Data ManagementServer is configured to use either a local Microsoft SQL Express database or an existing SQL Server instance. To change the database settings, use GdmsmsConfig.exe to change them using a graphical user interface (see chapter 17.1.8). GdmsmsConfig.exe saves settings automatically in Gdmsms.exe.config, where they can optionally be edited directly.

Setting	Default	Description
DbServer	.\GDATASQLSRV2K8	The name and instance of the database server, formatted as Server\Instance. Corresponds to the SQL Server setting in GdmsmsConfig.exe.
DbDatabase	GDATA_AntiVirus_ManagementServer	The name of the database. Corresponds to the Database setting in GdmsmsConfig.exe.
DbIntegratedSecurity	True	Use Windows authentication (True) or SQL Server authentication (False) to log on to the database. Corresponds to the Logon Authentication setting in GdmsmsConfig.exe.
DbDatabaseUser	<empty>	SQL Server authentication username. Empty when using Windows authentication. Corresponds to the Logon User name setting in GdmsmsConfig.exe.
DbDatabasePassword	<empty>	SQL Server authentication password (hashed). Empty when using Windows authentication. Corresponds to the Logon Password setting in GdmsmsConfig.exe.
DbImpersonateUser	<empty>	Defines the Windows user account which is used to log in to the database. Empty by default, SYSTEM account is used. Corresponds to the Logon User name setting in GdmsmsConfig.exe.
DbImpersonatePassword	<empty>	Defines the Windows user password which is used to log in to the database (hashed). Empty by default. Corresponds to the Logon Password setting in GdmsmsConfig.exe.
PerformStartupDbCheckAndRepair	True	Remove redundant database entries when the ManagementServer service is started.

17.1.2. Ports

During the installation of G Data ManagementServer, the server will assume the default values for

the TCP ports with which it connects to G Data Administrator and mobile clients (see chapter 2.2). The port numbers have been chosen to minimize impact on existing software. However, if there happens to be a port conflict, the port assignments can be changed in Gdmms.exe.config.

When changing the value for ClientHttpPort or ClientHttpsPort, the HTTPS security configuration for the port has to be reinitialized before restarting the ManagementServer service. To do so, open a command prompt with administrative privileges and run *C:\Program Files\G DATA\G DATA AntiVirus ManagementServer\gdmmsconfig.exe /installcert*. After changing the value for AdminPort, when logging on to G Data Administrator, the server port has to be explicitly defined in the format *servername:port*.

Setting	Default	Description
AdminPort	0	Port for TCP communication with G Data Administrator. Enter any port number. Value 0 sets the port to the standard number of 7182.
ClientHttpsPort	0	Port for TCP communication with mobile clients. ClientHttpsPort should not be altered, as mobile clients do not accept an alternative port. Enter any port number. Value 0 sets the port to the standard number of 7183.
ClientHttpPort	0	Port for TCP communication with mobile clients (distribution of mobile client installation files). Enter any port number. Value 0 sets the port to the standard number of 7184.

17.1.3. Server

In addition to the database and port settings, other ManagementServer functions can also be tweaked using Gdmms.exe.config. When changing server configuration settings, be aware that they often require changes in other contexts (e.g. G Data Security Client, Registry) to function correctly.

Setting	Default	Description
DisplayLicenseLimit	True	Display the number of permitted product licenses and the license expiration date in the License overview window's Extended view.
IsSecondaryMMS	False	Determines whether the server is a primary or secondary ManagementServer.
MainMms	<name>	Server name as defined during the installation of G Data ManagementServer, used to populate the server name for installation packages and remote client installations. Must be identical to HKEY_LOCAL_MACHINE\Software\G DATA\AVK ManagementServer\ComputerName.
MaxParallelClientInstallation	5	The maximum number of clients that can be installed remotely at the same time. Large numbers can lead to network congestion. Minimum 5, maximum 1000.
MaxSubnetUpdateThreads	300	The maximum number of subnet servers that can connect to the ManagementServer for updates or synchronization.
MaxUpdateThreads	300	The maximum number of clients that can connect to the ManagementServer for updates or synchronization. Ignored if the load limit has been enabled in G Data Administrator and a value for Simultaneous update download has been set.
DisableActiveDirectorySearch	False	Disable Active Directory synchronization.
EmailCodePage	<empty>	Code page for outgoing e-mails. When no code page is defined here, UTF-8 is used.

17.1.4. Folders

Folders where ManagementServer stores its data can be configured, although the default values should suffice for most scenarios. When redefining folders, existing files in those folders should be moved into the new folder before restarting the ManagementServer service.

When updating the value for UpdateDistributionFolder, the value for BasePath in IUpdateCfg.xml needs to be set to the same folder (see chapter 17.1.8).

Setting	Default	Description
BackupFolders	<folder>	Backups are stored by default in C:\ProgramData\G DATA\AntiVirus ManagementServer\Backup. Ignored if Server backup paths have been defined in G Data Administrator (see chapter 12.1).
DBBackupFolder	<empty>	The last used database backup folder in GdmmsConfig.exe. The folder setting should only be changed through the interface of GdmmsConfig.exe.
InternetUpdatePgmFolder	<folder>	The folder where the Internet Update component (IUpdate.exe) is stored, by default C:\Program Files (x86)\G DATA\G DATA AntiVirus ManagementServer. This value should not be changed.
PatchFilesFolder	<folder>	Patch management files are stored by default in C:\ProgramData\G DATA\AntiVirus ManagementServer\Patches.
QuarantineFolder	<folder>	Quarantined files are stored by default in C:\ProgramData\G DATA\AntiVirus ManagementServer\Quarantine.
UpdateDistributionFolder	<folder>	Virus signature updates and program file updates are stored by default in C:\ProgramData\G DATA\AntiVirus ManagementServer\Updates.

17.1.5. Peer-to-peer update distribution

Using the INTERNET UPDATE window in G Data Administrator, peer-to-peer update distribution can be enabled to reduce server load caused by update distribution (see chapter 7.3). Peer-to-peer update distribution can be used without additional configuration, but its settings can still be tailored to network circumstances, using Gdmms.exe.config. For example, "vicinity" can be defined using several parameters.

Setting	Default	Description
P2PClientAbandoned ConnectionThresholdMin	1	Maximum amount of inactivity time before a connection is considered abandoned (client side).
P2PClientMaxServedPeers	5	Maximum number of clients simultaneously served by one peer.
P2PConnectRetries	3	Maximum number of connection retries before a client forwards a connection request to another peer.
P2PConsiderClientsDHCP	False	Consider clients obtaining their dynamic IP addresses from the same DHCP server to be in each other's vicinity.
P2PConsiderClientsDomain	True	Consider clients from the same domain to be in each other's vicinity.
P2PConsiderClientsGateway	False	Consider clients using the same default gateway to be in each other's vicinity.
P2PConsiderClientsLastAccess	True	Consider the client's last access during determining a client as a source of an update.

P2PConsiderClientsOnBattery	False	Exclude clients running on battery power as source for updates.
P2PConsiderClientsSubnet	False	Consider clients from the same subnet to be in each other's vicinity.
P2PDisablePGMUpdate Distribution	False	Enable or disable peer-to-peer distribution of program file updates. (If disabled, only signatures updates will be distributed peer-to-peer).
P2PMaxNumberOfHops	1	Maximum number of hops between two clients in order to consider them as in each other's vicinity.
P2PMmsAbandoned ConnectionThresholdMin	1	Maximum amount of inactivity time before a connection is considered abandoned (ManagementServer side).
P2PMmsMaxServedPeers	50	Maximum number of simultaneous client connections served by the MMS (connections through which signature/program update is being downloaded)

When updating any of the peer-to-peer distribution settings, the changes have to be confirmed by editing a registry key before restarting the ManagementServer service. Using Registry Editor, remove the value DoNotConsiderP2PConfigToDB from the key HKEY_LOCAL_MACHINE\Software\G DATA\AVK ManagementServer (when using a 64-bit system, the key may be located under Software\Wow6432Node\G DATA). This forces ManagementServer to import the configuration values from Gdms.exe.config for peer-to-peer update distribution.

17.1.6. Staged software distribution

Staged software distribution is a very powerful tool, developed to allow administrators to distribute program updates across the network without having to worry about compatibility issues. Using mathematical principles, all network clients are divided into groups that get bigger each time a stage has been successfully completed. The system has been built to function without any configuration – simply enabling it is enough (see chapter 7.3.2). However, many parts of the calculation can be adjusted to perfectly fit the circumstances in every enterprise network.

For several settings, default values for the first six stages have been defined. When using more stages, additional values can be added to the configuration file. If there are more stages than values for a specific settings, the next values will be extrapolated based on the range.

Setting	Default	Description
SPUEnable	3	Bitmask for two settings in G Data Administrator's Server settings window: Distribute automatic software updates by stages (1) and Automatically allocate clients for the first stage (2).
SPUStopAbsolute	5,15,20,30,40,50	The number of corrupted clients per stage at which staged software distribution should be halted (absolute number). Used if smaller than SPUStopPromille.
SPUStopPromille	25,75,100,150,200,250	The number of corrupted clients per stage at which staged software distribution should be halted (per thousand). Used if smaller than SPUStopAbsolute.
SPUStepsTimespan	3.00:00:00 (3 days)	dd.hh:mm:ss. Amount of time until the release of the following stage. Corresponds to the Release next group after setting in G Data Administrator.

SPUStepsZombie Timespan	14.00:00:00 (14 days)	dd.hh:mm:ss. Clients that have not connected to the MMS for a certain amount of time will not be counted in the mathematical calculation that governs staged software distribution.
SPUFirstStepLimit	5,20	Minimum and maximum number of clients to be included in the first stage.
SPUTotalSteps	3	Number of stages. Corresponds to the NUMBER OF GROUPS setting in G Data Administrator.
SPUSyncTimespan	00:30:00 (30 minutes)	dd.hh:mm:ss. Synchronization of staged software distribution status between clients, subnet servers and MMS.
SPUMinClients	10	Minimum number of clients in the network required for staged software distribution to kick in.

17.1.7. Obsolete settings

Some of the settings in Gdmms.exe.config are no longer available in the latest version of G Data ManagementServer. For compatibility with upgrade installations, they are not removed from the file, but for the measures to be effective, an alternative configuration is required.

Setting	Alternative
LogFileFolder	Obsolete from version 9 onwards, as scan job logs are stored in the ManagementServer database (see chapter 9.1).
LogLevel, DebugLogDestination, DebugLogFilename	Obsolete from version 12 onwards, as debugging is configured using a separate configuration file (see chapter 17.3.3).
UseIntegratedDatabase	Obsolete from version 12 onwards, as the integrated database has been replaced by a SQL database (see chapter 17.1.1).

17.1.8. Server updates

G Data ManagementServer's program files can be updated through the Internet Update tool (see chapter 4.5). In addition, it can be used to perform manual updates of client program files and virus signatures. Some update parameters can be set in IUpdateCfg.xml.

Setting	Default	Description
BasePath	<folder>	Virus signature updates and program file updates are stored by default in C:\ProgramData\G DATA\AntiVirus ManagementServer\Updates.
UpdateCnt	5	The number of update rollbacks to be stored (see chapter 7.3.3). Corresponds to the value set on the Settings tab of the Server settings window of G Data Administrator.

When updating the value for BasePath, the value for UpdateDistributionFolder in Gdmms.exe.config needs to be set to the same folder (see chapter 17.1.4).

17.2. Configuration tools

In addition to the XML-based configuration files, several tools have been made available to fine tune G Data products. Some tools are installed with G Data ManagementServer, others can be downloaded from the G Data website if required. Like advanced configuration files, software tools need to be used with care. Before making any configuration changes, it is recommended to make a backup of all affected files and folders. Moreover, experimenting on a live network or network client

is not recommended. Before deploying any changes through a configuration tool, test its effects on a (virtual) test network or client.

17.2.1. Database maintenance

For server database backup and restore operations, the GdmmsConfig.exe tool is a very valuable tool (see chapter 4.6). Backups, however, are not its only purpose: GdmmsConfig.exe can also be used to carry out database maintenance. The interface shows the current data with which ManagementServer connects to its database. In its simplest configuration, this is a local SQL Server Express database with a local logon account and a default database name (GDATA_AntiVirus_ManagementServer). Depending on the type of deployment, however, database access can quickly become more complex.

Usually, database access details are automatically configured during the installation procedure for ManagementServer. GdmmsConfig.exe can be used if database parameters need to be changed after installation. It can become necessary to change the account with which ManagementServer logs in to the database. After entering the updated data into GdmmsConfig.exe, click TEST SETTINGS to make sure that the new data can indeed be used to successfully log in to the database. Another use case is database migration. For installations up to one thousand clients, a local SQL Server Express installation performs well (depending on the hardware configuration). If the ManagementServer hits one thousand clients, it may be necessary to move the database to its own server to ensure continuing performance. Using GdmmsConfig.exe, this process is relatively easy. Using the current database access details, create a database backup, then enter the new server's details and restore the database to the new server (see chapter 4.6). Finally, database versioning is also possible using GdmmsConfig.exe. Rather than just backing up and restoring the database, a new database can be created to restore the old data to, creating the possibility to experiment with different databases on the same server.

17.2.2. File system monitor activity

The file system monitor is one of the most thorough components of the G Data security solution. It detects a large amount of malware, but also needs considerable system resources to do so. When client performance is clearly impacted, the Monitor activity tool (MonActivityCS) can help find out if the file system monitor is having problems with a specific file. The tool can be downloaded from www.gdata.de/kundenservice/downloads/tools.html and can be run directly, without extracting it or running an installation wizard. The tool can be run in two modes: real time or hit list. The real time mode lists files as they are checked by the file system monitor. This is useful if the problem with the file system monitor can be easily reproduced. The hit list will show how often specific files were checked, which helps single out problematic files.

After choosing the mode, the main window of Monitor activity is shown. To start monitoring on the local machine, click Connect (the Computer text field can remain empty); click Stop to halt monitoring. The list will be populated as the tool is running. Click Refresh to manually update the list (in Hit list mode), Reset list to reset the list (in Real time mode) or Save to save the list of files as a text file. To immediately define a file as an exception for the file system monitor, click Define as exception.

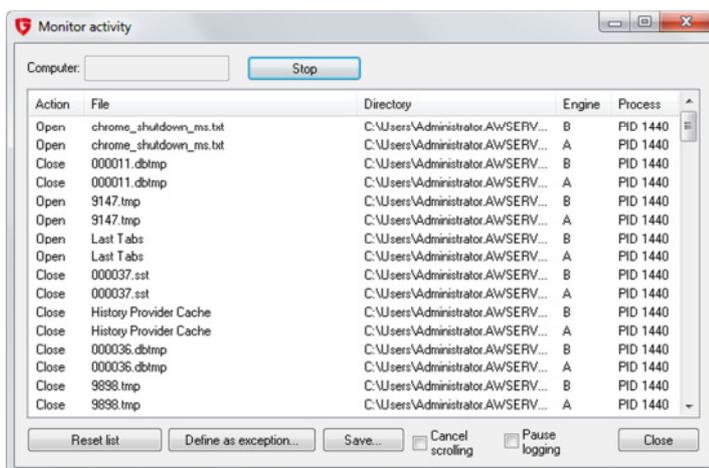


Image 76: Monitor activity

17.2.3. Quarantine

When one of the G Data security layers detects a threat, it is automatically taken care of. If the administrator has defined that files should be quarantined, they are renamed and moved to a secure folder. Using G Data Administrator’s REPORTS module, the files can then be analyzed, cleaned and moved back. To manage a client’s quarantine locally, the Quarantine tool can be used. It can be downloaded from www.gdata.de/kundenservice/downloads/tools.html (Quarantine generation 2011 and newer) and should be run on the affected client. The tool can also be used to check files that are located in the server’s quarantine folder, in case analysis through G Data Administrator is not possible.

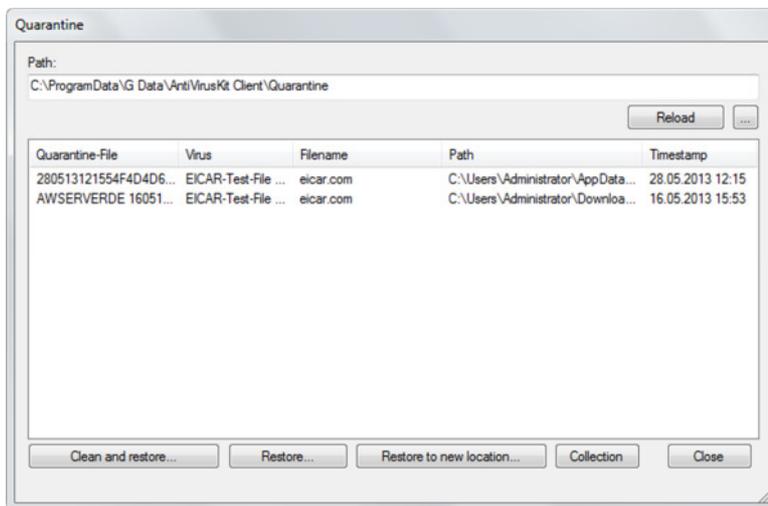


Image 77: Quarantine

Once it has started, the Quarantine tool immediately displays the local quarantine. It takes the default quarantine path and lists all files, along with some metadata. If it does not automatically select the correct folder, it can be picked manually. By default, client quarantine files are saved in C:\ProgramData\G Data\AntiVirusKit Client\Quarantine. When examining the server quarantine, the default folder is C:\ProgramData\G Data\AntiVirus ManagementServer\Quarantine, unless a different path has been specified in Gdmms.exe.config (see chapter 17.1.4).

When selecting a quarantined file, several options are available. The recommended option is CLEAN

AND RESTORE. This will attempt to remove the malware from the file and restore it to its original location. Other options are to restore the file to its original location, without cleaning it, and to restore the file to a new location without cleaning it. Both of these are not recommended: the malware will still be present in the file and poses a risk. If a file cannot be cleaned, it may have to be deleted. To do so, open the Quarantine folder in Windows Explorer and delete the associated .q file.

17.2.4. Uninstallation

G Data ManagementServer and G Data MailSecurity can be easily removed using the uninstallation wizard (CONTROL PANEL > ADD/REMOVE PROGRAMS). For other components, uninstallation can be initiated remotely. G Data Security Client, for example, can be uninstalled through G Data Administrator's CLIENTS module. The G Data Firewall can be removed from clients using the FIREWALL module. When using one or more subnet servers, the function `MANAGE SERVER(S)` can be used to initiate a remote uninstallation.

G Data ManagementServer can be locally uninstalled using a command prompt. This is especially useful for subnet servers that can no longer be accessed remotely. On the affected server, open a command prompt with administrator permissions and navigate to the installation folder of G Data ManagementServer (typically `C:\Program Files (x86)\G DATA\G DATA AntiVirus ManagementServer`). Run the command `cuninstmms /uninst`. This will shut down the ManagementServer background service and initiate the uninstallation wizard for G Data ManagementServer. After confirming whether database contents should be kept or removed, the wizard will uninstall G Data ManagementServer in the background.

G Data Security Client can be remotely uninstalled by selecting the appropriate client in G Data Administrator and selecting `UNINSTALL G DATA SECURITY CLIENT`. Like ManagementServer, G Data Security Client can also be uninstalled locally. Open a command prompt with administrator permissions and navigate to `C:\Program Files (x86)\G DATA\AVKClient`. Enter the command `unclient /AVKUninst` to start the procedure. The client should be restarted afterwards.

When uninstalling a subnet server or Security Client locally, related data in the ManagementServer's database are not automatically removed. Make sure to remove any inactive subnet servers using the `MANAGE SERVER(S)` function. Inactive clients can be pruned in the client management area.

17.2.5. Uninstallation cleanup

All G Data software components can be uninstalled using simple tools (see chapter 17.2.4). However, if the uninstallation procedure is not completely carried out, future installations may fail because of temporary registry keys or files. To make sure that there are no leftovers, the AV-Cleaner tool can be run. AV-Cleaner is available from www.gdatasoftware.co.uk/support/downloads/tools.html. It should be run directly on the affected client or server, after the regular uninstallation procedure has completed.

Upon execution, AV-Cleaner gathers information about files, services, registry keys and other objects that are related to G Data and displays them in a list. Files and keys can be deleted individually, or they can all be selected at once to clean up everything. After removal, the computer should be restarted and AV-Cleaner should be run again to make sure there are no traces left. Repeat the reboot and AV-Cleaner cleanup until no traces are left.

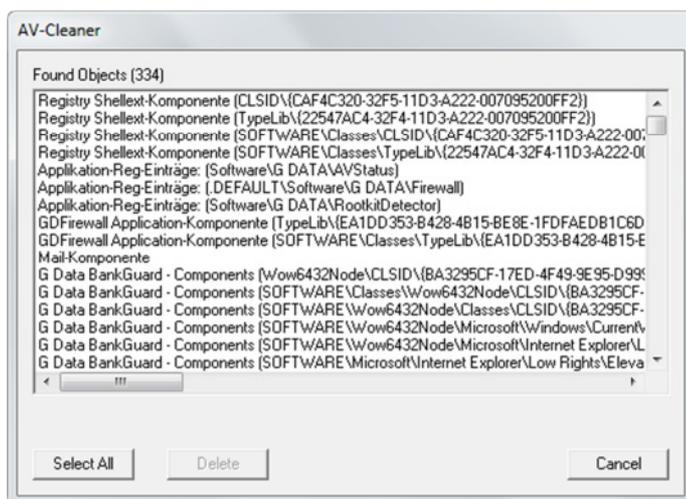


Image 78: AV-Cleaner

Note that AV-Cleaner completely removes all G Data software. This means that extra care should be taken when removing a G Data product from a computer on which other G Data products have been installed. For example, when removing G Data Security Client from a computer on which G Data ManagementServer has been installed as well, it is strongly recommended to use the regular uninstallation method to prevent any ManagementServer components from being removed by AV-Cleaner.

17.3. Debug logs

To assist in advanced configuration and troubleshooting, it can be helpful to examine debug logs. Many of the client-side security modules produce their own logs that can be read in G Data Administrator, such as virus scans, backup jobs or restore jobs. Additional debugging can be done using several configurable logs. Using those files, the inner workings of most security modules can be analyzed. If any part of the software is acting unexpectedly the debugging log will provide valuable insights. Even if direct analysis is not possible, the log files can help our support department investigate the problem. Generating debug logs for ManagementServer itself can also be useful in several cases. If a client does not receive its updated settings, a subnet server causes problems or the ManagementServer is using extraordinarily large amounts of RAM or CPU cycles, debug output may give an indication of the type of problems occurring.

There are several ways of gathering log files, depending on the module that should be monitored. For most of the security modules, the DebugView method lets administrators easily gather logs. If the ManagementServer installation should be analyzed, the Debug.config method provides logs. To debug the installation wizard of ManagementServer, setup parameters can be set.

17.3.1. Installation wizard debugging

The installation wizard of G Data ManagementServer provides on-screen guidance to help administrators install the software without problems. If, after the installation wizard concludes, some components of ManagementServer cause problems or do not run at all, it can be useful to check if the wizard properly installs all components, or if exceptions are occurring. The installation wizard can be configured to generate a debug log, which lists all actions that it takes and any errors. Open the Registry editor and navigate to the key HKEY_LOCAL_MACHINE\Software\G DATA\ (on 64-bit

systems: HKEY_LOCAL_MACHINE\Software\Wow6432Node\G DATA\) and add the key *GDSetup*, if it does not exist already. Within the new key, add a DWORD value called *DebugLevel* with value 5. Open a command prompt and navigate to the folder that contains the ManagementServer setup file (when installing from the installation medium, typically D:\Setup\ManagementServer). Type the following command to execute the installation wizard with the option to generate an extensive debug log: `setup.exe /params=[/v"/!*"v c:\debug.log"]`. This will log all setup actions verbosely to c:\debug.log, which can then be analyzed afterwards. To debug an installation of G Data MailSecurity Exchange plugin, open registry key HKEY_LOCAL_MACHINE\Software\G DATA\Exchange and add the DWORD value *DebugLevel* with value 5. Using the DebugView tool, the installation's debug output can be captured (see chapter 17.3.2).

17.3.2. Client-side debugging

Client-side debugging can be carried out using the DebugView tool, developed by Sysinternals. It intercepts debug calls for all Win32 software, which allows administrators to gather debug logs even without an active debugger. In combination with setting a registry key for the module that should be debugged, DebugView generates extensive logs. The same method is used to debug MailSecurity MailGateway.

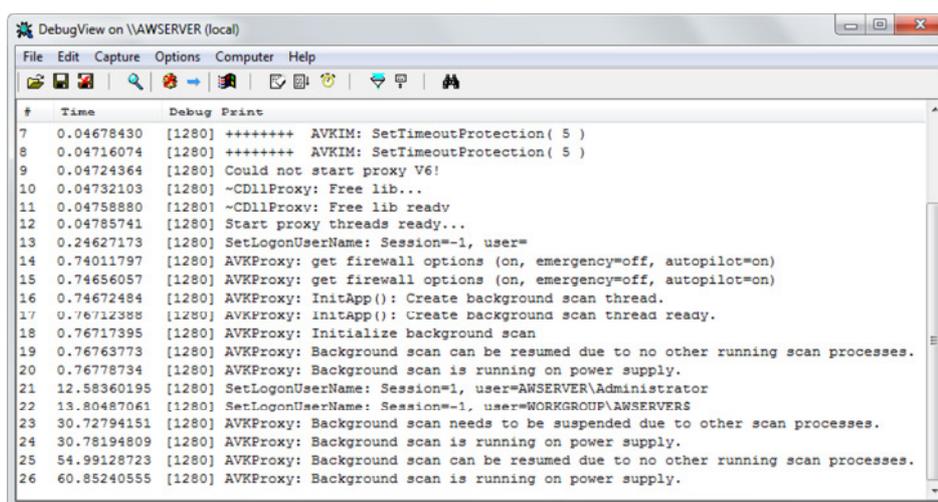


Image 79: DebugView

Using the registry editor, add a DWORD value called *DebugLevel* with the data 5 to the following keys to allow debugging for the respective modules:

Module	Key ¹⁸
MailSecurity MailGateway	HKEY_LOCAL_MACHINE\Software\G DATA\AVKSmtp
Security Client (general)	HKEY_LOCAL_MACHINE\Software\G DATA\AVKClient
Traffic scans	HKEY_LOCAL_MACHINE\Software\G DATA\AVKProxy
Updates	HKEY_LOCAL_MACHINE\Software\G DATA\INetUpdate
Virus scans	HKEY_LOCAL_MACHINE\Software\G DATA\AVKScanP

After the registry change, restart the PC. Download DebugView from <http://technet.microsoft.com/en-us/sysinternals/bb896647>. After downloading the file, extract it

¹⁸ On 64-bit systems, the G DATA key is located in HKEY_LOCAL_MACHINE\Software\Wow6432Node\G DATA.

and run DebugView.exe with administrator permissions. Under CAPTURE, enable CAPTURE GLOBAL WIN32. While DebugView is running, it will collect debug output from the modules for which the DebugLevel registry key was set. The DebugView window will show all output, which can be saved to a text file. When using DebugView to investigate an error, run it until the error can be reproduced, then save the log. The log can be used for further manual analysis, or be sent to our support representatives.

An exception to the general debug instructions is Device control (one of the PolicyManager modules). Because Device control requires deep level access to Windows' device settings, additional registry settings are required to enable debugging. In the Registry Editor, navigate to the key HKEY_LOCAL_MACHINE\System\CurrentControlSet\Control\GDDevCtrl, creating the key if it does not exist already. Create a DWORD value called *DebugLevel* with the data 5. Navigate to HKEY_LOCAL_MACHINE\System\CurrentControlSet\Control\Session Manager\Debug Print Filter, creating the key if it does not exist already. Add two new DWORD values: *Default* with data *ffffff* (hexadecimal), and *IHVDRIVER* with data *ffffff* (hexadecimal). Start DebugView as administrator, enable FORCE CARRIAGE RETURNS under OPTIONS, and CAPTURE KERNEL and ENABLE VERBOSE KERNEL OUTPUT under CAPTURE. Restart the client and restart DebugView. This enables debugging for Device control, and will provide relevant logs when the problem is reproduced.

17.3.3. Server-side debugging

To configure debugging for the various components of ManagementServer, the separate configuration file Debug.config is used. It is not included with the default installation of ManagementServer, but can be easily created in a text editor such as Notepad. The file should be saved in the installation folder of G Data ManagementServer, typically C:\Program Files (x86)\G DATA\G DATA AntiVirus ManagementServer. For ManagementServer to pick up any changes in Debug.config, the ManagementServer background service needs to be restarted. Make sure to do this each time changes to Debug.config are saved.

```
<?xml version="1.0" encoding="utf-8" ?>
<debug>
  <output file="true" fileName="c:\temp\debug.log" console="false" />
  <categories>
    <category name="Root" level="4" />
    <category name="Database" level="4" />
    <category name="PatchManagement" level="2" />
    <category name="Sync" level="2" />
    <category name="WCF" level="4" />
  </categories>
</debug>
```

Image 80: Debug.config

Debug.config follows an XML structure, with several parameters that can be customized. The debug output can be read using DebugView (see chapter 17.3.2) – optionally, file output can be enabled. When channeling debug output to a file, debugging should only be enabled while troubleshooting. As soon as the appropriate logs have been collected, disable debugging to prevent the log file from quickly filling up the hard drive.

Output	Description
file	True or False. Save debug output to a file.
fileName	Output file including path, for example C:\MMSDebug.log.

console True or False. Internal debug function.

In the Categories section, several types of actions can be logged. Each category has a name attribute, which defines which type of action is logged, and a level attribute, which defines how detailed the logs will be. The level can be any number from 0 to 6, each level including information from all previous levels as well. Level 0 does not log anything; level 1 logs only errors; level 2 logs warnings; level 3 logs information; level 4 enables verbose information logging; level 5 logs everything in debug mode; level 6 logs everything in trace mode. For troubleshooting purposes, level 4 is recommended.

Category	Description
Root	The Root category sets the general debug level. No other category can exceed the Root category; if the Root category is set to 0, no debug output will be generated.
Database	SQL transactions between ManagementServer and its database.
PatchManagement	PatchManager events.
Sync	Database synchronization.
WCF	Communication between Administrator and ManagementServer.

To enable debugging for MailSecurity MailGateway, see chapter 17.3.2.

Acronyms

AD	Active Directory
API	Application programming interface
APK	Application package file (Android)
AV	Antivirus
BCC	Blind carbon copy
BIOS	Basic input/output system
CC	Carbon copy
CPU	Central processing unit
CVE	Common Vulnerabilities and Exposures
DMZ	Demilitarized zone (network)
DNS	Domain name system
ERP	Enterprise resource planning
EULA	End user license agreement
FTP	File transfer protocol
GPS	Global positioning system
HIPS	Host-based intrusion prevention system
HTTP	Hypertext transfer protocol
IIS	Microsoft Internet Information Services
IM	Instant messaging
IMAP	Internet message access protocol
IT	Information technology
JS	JavaScript
MD5	Message-Digest Algorithm 5
MMS	G Data ManagementServer
MX	Mail exchanger (record)
OU	Organizational unit
PC	Personal computer
POP3	Post office protocol version 3
R&D	Research & development
RAM	Random access memory
RBL	Real-time blacklist
SD	Secure Digital
SIM	Subscriber identity module
SMB	Small and medium-sized businesses
SMS	Short message service
SMTP	Simple mail transfer protocol
SP	Service pack
SQL	Structured query language
SSID	Service set identification
SSL	Secure socket layer
TCP/IP	Transmission control protocol/Internet protocol
UAC	User account control
UDP	User datagram protocol
UNC	Uniform naming convention
URL	Uniform resource locator
UTF-8	Universal Character Set transformation format – 8-bit
VBS	Visual Basic script
WAN	Wide area network
WCF	Windows Communication Foundation
WLAN	Wireless local area network
XML	Extensible markup language

Index

Active Directory.....	51
administration.....	25
alarms.....	48
application filtering.....	101
apps.....	82
autorun.....	65
backups.....	86
BankGuard.....	59
client role.....	7
clients.....	50
deployment.....	32
installation package.....	34
local installation.....	34
remote installation.....	33
troubleshooting.....	35
Linux clients.....	37
management.....	50
mobile clients.....	<i>See mobile device management</i>
remove.....	55
troubleshooting.....	35
configuration.....	25
browser.....	40
configuration files.....	137
configuration tools.....	142
desktop application.....	39
MasterAdmin.....	43
mobile.....	42
Dashboard.....	44
deployment.....	16, 17
device control.....	103
firewall.....	94
rule sets.....	96
groups.....	50
heuristics.....	62
idle scan.....	67
internet usage time.....	106
licensing.....	15
load limit.....	<i>See performance</i>
logs.....	
client.....	147
firewall.....	99
installation.....	146
server.....	148
virus scan.....	<i>See malware infection</i>
mail groups.....	28
MailSecurity.....	26, 116
Gateway.....	118
Microsoft Exchange.....	26, 116
malware analysis.....	78
malware infection.....	75
ManagedService.....	21
MasterAdmin.....	43
mitigation.....	75
mobile device management.....	80
apps.....	82
deployment.....	36
policies.....	81
protection.....	80
network access monitoring.....	62
network diagram.....	6
network zone.....	6
online registration.....	30
patch management.....	108
performance.....	55
PolicyManager.....	101
port numbers.....	14
product activation.....	25
quarantine.....	76
local analysis.....	144
ReportManager.....	48
reports.....	46
scan jobs.....	68
secondary ManagementServer.....	19
security components.....	11
server database.....	24
backup and restore.....	30
migration.....	143
Server setup wizard.....	28
solutions.....	12
spam filter.....	
MailSecurity.....	130
Security Client.....	59
statistics.....	44
subnet server.....	20
synchronization.....	29
system requirements.....	13
traffic scans.....	58
updates.....	51
clients.....	52
distribution.....	53
peer-to-peer.....	54
schedule.....	52
staged distribution.....	53
offline update.....	52
rollbacks.....	54
server.....	29
upgrade path.....	27
vulnerability.....	<i>See patch management</i>
web filter.....	105