



PRT-GX-SRVR

Protege GX Large System Guide



The specifications and descriptions of products and services contained in this document were correct at the time of printing. Integrated Control Technology Limited reserves the right to change specifications or withdraw products without notice. No part of this document may be reproduced, photocopied, or transmitted in any form or by any means (electronic or mechanical), for any purpose, without the express written permission of Integrated Control Technology Limited. Designed and manufactured by Integrated Control Technology Limited, Protege® and the Protege® Logo are registered trademarks of Integrated Control Technology Limited. All other brand or product names are trademarks or registered trademarks of their respective holders.

Copyright © Integrated Control Technology Limited 2003-2026. All rights reserved.

Last Published: 28-Aug-26 11:08 AM

Contents

Introduction	4
Additional Documentation	4
System Architecture	5
Server Specifications	6
Databases	6
Controller Downloads	7
Events and Statuses	8
Silent Client Installation	8
Networking and Cybersecurity	8
Server Failover	9
Controller Capacity	11
Partitioning the System	12
Cross Controller Operations	13
Connecting more than 64 Controllers	13
Event Management	14
Disabling Unneeded Events	14
Backing Up and Purging the Events Database	14
Maximum Event Database ID	14
Integrations	16
Programming	17

Introduction

Protege GX often provides access control, intruder detection and user management for large enterprise customers such as airports, hospitals, high-rise buildings and universities. These large systems often encounter challenges in computing capacity, network architecture, download speed and event storage as the size of the system increases. In addition, enterprise customers frequently have high cybersecurity standards, complex networks and important integrations.

This document outlines the best practices for deploying Protege GX systems at enterprise scale. We recommend that you review these guidelines before deploying systems that contain:

- 200+ controllers, or
- 20,000+ users, or
- 400+ doors

Even if you are not deploying Protege GX at this scale, the recommendations and tips in this document may be helpful for optimizing your systems.

If you need more help with designing or deploying a Protege GX system, get in touch with ICT Technical Support ([Contact Us](#)).

Additional Documentation

This document is supplementary to the standard manuals for the Protege GX system. Ensure that you follow the instructions in the [Protege GX Installation Manual](#), as well as any application notes and other documents relevant to the system.

System Architecture

This section includes recommendations and options for planning large Protege GX systems.

Every Protege GX system has the following components:

- **Protege GX Server:** Hosts the Protege GX services.
- **SQL Server Instance:** Hosts the programming and event databases. May be on the same machine as the Protege GX server or a different machine.
- **Protege GX Thick Client:** Software for programming, viewing and reporting on the system. Installed on the server and workstations.
- **Protege GX Controllers:** Intelligent devices that manage access, security and automation on site.


Large systems typically have additional components:

- **Protege GX Web App:** Browser-based interface, enabling operators to complete everyday tasks from the web browser. It can be easier for large organizations to use a browser-based interface rather than installing software on every device that needs access.

Existing systems may use the older Protege GX Web Client.

- **Protege GX Enterprise Download Server:** Improves download speed and capacity for the site. Hosted on a virtual machine (VM), typically deployed to a separate computer.
- **Protege GX Extended Services:** Additional services that extend the Protege GX system, typically installed alongside the web app. These include the extended event service, which improves event reliability.
- Servers for third-party software that is integrated with Protege GX.

Overview of Protege GX System Architecture



Server Specifications

It is vital that you install Protege GX on a sufficiently powerful computer for smooth operation. Large systems may require higher server specifications than the minimum requirements stated in the Protege GX Installation Manual.

We recommend the following specifications for large systems:

- **CPU:** At minimum 64-bit, 6 cores, 12 threads, 2.5 GHz.

We recommend selecting the newest available CPU generation at time of purchase.

- **RAM:** 32-64 GB
- **Storage:** 500 GB - 1 TB (the capacity required depends on the amount of data that will be stored in the Protege GX databases)

We recommend using an SSD instead of an HDD, as they provide better performance.

- **Ethernet Connection:** 100 Mb/1 Gb

The Protege GX server may be installed on a virtual machine, as long as the VM has sufficient dedicated computing resources available. This can be hosted locally or by a cloud provider such as Microsoft Azure or Amazon Web Services.

Databases

The Protege GX databases are hosted in SQL Server.

Database Location

The databases can be hosted on the same machine as the Protege GX server or a different machine.

Hosting the databases on a different machine can be good for performance in large systems, as long as the network connection between Protege GX server and the database server is fast and robust. This also allows you to make use of existing IT infrastructure, such as sharing SQL instances between multiple applications.

When you host the databases on a remote computer, you must ensure that the Protege GX service account has sufficient access to the databases. See [Application Note 365: Restricting Protege GX Service Permissions](#) for the required permissions.

SQL Server Editions

The **Standard** edition of SQL Server provides enough storage and computing capacity for large Protege GX systems. This edition does have some limits on computing capacity, but these should not affect the performance of Protege GX (see [SQL Server Scale Limits](#) in the Microsoft Help).

The Enterprise edition should not be necessary to run the Protege GX server. However, if the organization is using the same SQL Server instance for multiple services, they may wish to use the Enterprise edition to ensure that there are no performance limitations.

SQL Server Express has a 10GB storage limit per database. Do not use it for large sites.

Database Backups

You must back up your databases at least once every 24 hours. The standard backup settings in the Protege GX software are:

- **ProtegeGX (programming) database:** Full database backup once per day. One week of backups is stored.
- **ProtegeGXEvents database:** Differential database backup once per day. The new backup is appended to the backup from the previous week.

For more information about event database backups, see [Application Note 279: Purging and Restoring the Protege GX Events Database](#).

If your site requires different backup settings (such as more transaction log backups, more frequent backups, or backups stored for longer), you must configure them in SQL Server. For instructions, see [Use the Maintenance Plan Wizard](#) in the Microsoft documentation.

Backup Storage

Backups can occupy a considerable amount of space on the database server, which can potentially cause the server to run out of disk space. We recommend implementing a disk space management strategy, such as:

- Automatically deleting old backup files
- Automatically moving old backup files to an external drive
- Setting up a low disk space alert so that an administrator can move or delete backup files

To ensure that you can recover from a server loss, we recommend that you store copies of database backups in a secure offsite location.

Recovery Model

The recovery model for SQL Server databases is set to *Full* by default. This means that every change to the programming and event databases is logged in the transaction log, allowing you to recover the databases to any point in time in case of a failure. However, saving such a large number of transactions can impact the performance of the system and quickly fill up disk space on the server.

For most sites we recommend changing the recovery model to *Simple* for both databases to reduce disk usage and performance impacts. With this recovery model the server does not log all changes, so databases can only be restored to the last full or differential backup.

However, sites with strict compliance requirements may need full transaction logging, so consult with the IT administrator.

For instructions on changing the recovery model, see the [Protege GX Installation Manual](#). For more information, see [Recovery Models](#) in the Microsoft documentation.

Controller Downloads

The standard Protege GX Download Service works by checking each controller in the system in order. When it finds a controller with pending changes, it creates a full download package and sends it to that controller. Then it continues on to the next controller.

As the number of controllers in the system increases, it takes longer for any one controller to receive programming changes. This can become very slow in large systems, causing considerable operational issues.

When planning systems with 20 or more controllers, we recommend you implement one of the solutions below to improve download times.

Protege GX Enterprise Download Server

The enterprise download server is the **recommended solution** for large systems. It is designed to handle hundreds of controllers, thousands of doors and tens of thousands of users, and offers considerable improvements in speed and performance over the standard download server. Most programming changes are downloaded to all controllers within seconds.

The enterprise download server is deployed on a dedicated virtual machine. It can require significant computing resources, so it typically needs to be deployed on a separate machine from the Protege GX server. It can also be deployed to existing VMware ESXi infrastructure.

Contact ICT Technical Support if you're interested in using this feature.

Protege GX Single Record Download Service

The single record download service greatly improves download times for users, access levels and schedules only. This mitigates issues with the most common operations, such as granting access to new users.

This is a separate service that is typically installed on the same machine as the Protege GX server. It is a relatively lightweight option, so may be useful in situations where you do not have another server available.

For more information, see [*Application Note 309: Single Record Downloads in Protege GX.*](#)

Secondary Download Servers

You can install additional standard download servers on other computers to increase the download throughput. Each download server manages a specific set of controllers (e.g. two download servers, each managing half of the controllers).

It may be helpful to reserve one download server to use for priority downloads (sometimes called a 'maintenance' server). This server does not manage routine downloads and has no controllers permanently assigned to it. If you need to download a change quickly, you can temporarily assign a controller to this server to immediately trigger a full download.

For more information, see [*Application Note 290: Setting Up a Secondary Protege GX Download Server.*](#)

Events and Statuses

In a default system, events and status updates from controllers are handled by the Protege GX Event Service. However, for large systems we recommend installing the **Protege GX Extended Event Service** alongside the existing services.

The extended event service takes over the function of receiving events from controllers. Compared to the standard event service, it is faster, more scalable and more reliable, significantly reducing the risk of lost events.

There are two ways to install the extended event service:

- When you install the Protege GX Web App, the extended event service is also installed automatically. See the [*Protege GX Web App Installation Manual*](#) for instructions on setting it up.
- Alternatively, you can install the extended event service independently, without the web app. For instructions, see [*Application Note 363: Protege GX Extended Event Service*](#) (available on request from ICT Technical Support).

See [*Event Management*](#) below for tips on managing event storage.

Silent Client Installation

You can install, upgrade and uninstall the thick client via the command line, making it easier to deploy Protege GX to workstations across the organization. For more information and instructions, see [*Application Note 167: Protege GX Silent Client Installation.*](#)

Networking and Cybersecurity

Enterprise organizations often have complex network setups and high cybersecurity standards. Always engage with the IT administrator early to ensure that you set up the Protege GX system correctly.

We strongly recommend using a dedicated, segmented network for the security system. Only authorized devices and traffic should be allowed to enter this network.

Use the following two documents for guidance when setting up the system:

- [*Protege GX Network Administrator Guide*](#)
- [*Protege GX System Hardening Guide*](#)

Server Failover

When the Protege GX server goes down, controller-based operations such as access control, intruder detection and central station alarm reporting will continue to run as normal. Each controller can store up to **50,000 events** while it is offline from the server, after which it must start deleting the oldest events.

However, functionality provided by the software, such as updating configuration, monitoring and reporting, will not be available for as long as the server stays down.

On large or high-security sites, it is important to implement server redundancy and a failover process to ensure that the server can recover quickly after a failure. Below are a few options for configuring failover.

ICT cannot help with the setup requirements for third-party products. Contact the third-party provider for assistance.

Continuous Availability

To ensure continuous availability of the server, use a downtime prevention product such as [*Stratus everRun*](#). These products run redundant hardware in parallel, so a hardware failure causes little to no downtime.

Cold Standby

Alternatively, you can maintain a second, redundant Protege GX server. If the primary server fails or needs to be taken offline, you must manually switch the system over to the secondary server.

We recommend practicing the cold standby failover process on a regular basis to ensure that everything works as expected and the procedure can be completed quickly.

A few considerations for using this method:

- **License:** The Protege GX server license is tied to the hardware profile. Each server will need a separate Protege GX server license with all licensed features.
- **Services:** The Protege GX services of the secondary server should be stopped while the primary server is running. When the secondary server is needed, start the services.
- **Ports:** The secondary server must use the same event, download and control ports as the primary server.
- **Firewall:** The secondary server must have the firewall configured to allow connections over these ports.
- **Databases:** You must keep the secondary server up to date with the latest configuration and events. Use an automated process to restore databases from the primary server to the secondary server on a regular basis. When the secondary server comes online, it will contain all data up to its last update. Any changes made after the most recent update are lost.
- **Encryption:** If the primary server uses transparent data encryption (TDE) or column encryption in SQL Server, the secondary server must have the same encryption certificates installed.

Controller encryption keys are stored in the database, so the controllers will be able to communicate with the secondary server as long as the databases are up to date.

- **Event Server:** In each controller's web interface, set **Event Server 2** to the secondary server's IP address. When you start up the event service on the secondary server, controllers will automatically begin sending events to it.
- **Computer Name:** The download and event services rely on the computer name being configured correctly in Protege GX. When you bring the secondary server online, rename the secondary server to match the primary.

- **Health Status:** If the primary server fails partway through a download, the secondary server completes that download when it starts up. This can stop some programmable functions and services. After bringing the secondary server online, monitor the controllers' health statuses to see whether any functions or services have stopped, then start them again if necessary.

Controller Capacity

Maximum Capacities

The table below shows the **maximum** capacity for records that can be saved on the controller. Most of these capacities are set by the controller's internal database structure. However, the actual performance of the controller may degrade before this capacity is reached.

Record Type	Limit per Controller
Users	10,000
Events	50,000
Expander Modules	
Expander Modules (Total across all module types)	250
Analog Expanders	32
Input Expanders	248
Keypads	200
Output Expanders	32
Reader Expanders	64
Smart Readers	248
Other Records	
Data Values	248
Doors	248
Elevator Cars	8
Floors	128
Programmable Functions	248
Services	16

Recommendations

- We recommend connecting no more than **150 expander modules** to each controller. Connecting more modules can cause performance degradation, such as the controller taking longer to respond to access requests.

Partitioning the System

The Protege GX system can be divided into multiple partitions, which may represent locations, customer accounts, security levels and other divisions. It is important to plan the partitioning of the system before you begin. This ensures that the organization can effectively manage operator access and filter large numbers of records.

Protege GX offers two levels of partition: sites and record groups.

Sites

Sites are the major divisions of the Protege GX system. Most Protege GX servers only contain one site, but there are some scenarios where a server may have multiple sites.

Sites have the following features:

- All sites share global settings, but can have different site defaults.
- Sites share operators. Each operator may be permitted to access only one site or multiple sites, and they may have different permissions in each site.
- Sites do not share users, access levels, physical hardware or other operational functions. It is not possible to move records between different sites. For example, if a cardholder needs access to multiple sites, you must create a separate user record for the same person on each site.
- Each Protege GX site can include multiple physical locations.

Use sites when you need to manage more than one Protege GX installation on the same server. For example, if you are running a hosted service model with several customers on the same server, each customer should have their own site.

Record Groups

Record groups are flexible categories that can represent any kind of partition within a site. They are primarily used to manage operator access, although they can also be helpful for grouping and filtering records.

Record groups have the following features:

- Most records within a site can be assigned a record group, including users, access levels, doors, groups and so on. Physical hardware (e.g. inputs and outputs) uses the record group assigned to the controller.
- Each record can only be assigned one record group.
- Each record group may have a parent record group, allowing you to create a hierarchical structure.
- Each operator can access one or multiple record groups and will not be able to view any records outside of their groups. The operator must have the same permissions in each record group.
- You can filter records and events by the record group. Use the tree view to group users and access levels by record groups.

Record groups are commonly used to represent:

- Physical locations, such as a single branch location.
- Geographical areas such as cities or regions, using a hierarchical structure.
- Individual companies within a larger complex or multistory building.
- High-security zones such as server rooms and vaults.

For more information about managing operator access with record groups, see [*Application Note 191: Programming Operator Roles in Protege GX*](#).

Cross Controller Operations

Cross controller operations enable a group of controllers to communicate with each other, sharing hardware resources and tracking statuses across the system. This has a wide variety of uses, such as site-wide lockdowns, shared reporting and global antipassback.

When planning a large site, you should give consideration to the **cross controller groups** that you need to create in **Sites | Cross controller groups**. Cross controller groups have a *maximum size* of 64 controllers.

However, cross controller communication has an impact on the performance of controllers, as each controller has to keep track of inputs, outputs, doors and areas on other controllers. **We recommend that each cross controller group contains 48 controllers or fewer.**

For more information, see [*Application Note 180: Protege GX Cross Controller Operations*](#).

Connecting more than 64 Controllers

Each cross controller group can only include up to 64 controllers, and it is not possible to add a controller to more than one group. However, sometimes it is necessary to connect more than 64 controllers together to achieve a specific requirement for a large site.

For example, the following scenarios might require many controllers to coordinate with each other:

- When anyone presses a panic button, lock down the entire university campus.
- When a fire alarm is activated, unlock every door in the high-rise building.
- When an alarm occurs anywhere in the airport, activate a siren in the control center.

In these instances, you must connect multiple cross controller groups together. There are two methods to achieve this:

1. **I/O Wiring:** Physically wire an output from Cross Controller Group A to an input from Cross Controller Group B. When that output is activated (e.g. by an alarm or programmable function in Cross Controller Group A), the input is opened. That input can then trigger an alarm or activate an output in Cross Controller Group B.
2. **Link Me Service:** The Link Me service enables you to mirror one or more outputs between two controllers. When the primary controller activates or deactivates specific outputs, the secondary controller updates its own outputs to match. This allows Cross Controller Group B to control a programmable function, schedule or virtual input based on the status of the outputs in Cross Controller Group A.

See the [*Protege GX Operator Reference Manual*](#) for more information about the Link Me Service.

The best method for each site will depend on the number of outputs that need to be shared, the physical wiring situation and the network configuration.

We recommend keeping the controller group size to 48 controllers or fewer. It is better to create multiple smaller cross controller groups than a few large groups.

Event Management

The Protege GX system generates large numbers of events, increasing with the size of the site and how many people pass through every day. It is necessary to strike a balance between storing enough events for auditing and incident response without overwhelming the disk space of the database server.

Disabling Unneeded Events

Some events generated by the system are useful for initial installation and commissioning of the site, but not necessary in day-to-day operation. We recommend that you disable the following events after commissioning.

Disabled events are not generated by the controller at all. There is no way to recover any events that have been disabled.

Input Events

Many inputs, such as motion sensors and door inputs, are triggered frequently throughout the day. Most of these openings and closings are incidental and do not need to be recorded unless they cause an alarm, so it is common to suppress the 'Input Opened' and 'Input Closed' events. This will not prevent secondary effects such as alarms, reports and door responses.

To suppress input events, navigate to **Programming | Inputs | Options** and disable **Log to event buffer**.

Output Events

You can also disable events for outputs that do not need to be monitored, such as lights and door outputs.

To suppress output events, navigate to **Programming | Outputs | Options** and disable **Log output events**.

Door Events

It is typically not necessary to record openings and closings when a door is unlocked by schedule.

To suppress these events, navigate to **Programming | Doors | Options** and disable **Enable open/close events on schedule**.

Backing Up and Purging the Events Database

We recommend that you purge the events database regularly to prevent the storage from being filled. The frequency will depend on the number of events generated on site. Differential backups enable you to save the purged events and restore them if they need to be reviewed.

For more information and instructions, see [*Application Note 279: Purging and Restoring the Protege GX Events Database*](#).

Maximum Event Database ID

The events database has a maximum Event ID of 2147483647 (around 2.1 billion), which may be reached on very large and busy sites (200+ controllers) or those that have been running Protege GX for a long time. This is not affected by purging the events database.

We recommend archiving and replacing the events database before it reaches this limit. You can view the **Event ID** column in the *All Events* report or a status page (if this has been hidden, right click on any column header and select **Show column chooser** to retrieve it). When you observe new Event IDs exceeding one billion (ten digit numbers), **contact ICT Technical Support** for assistance with backing up the events database and creating a new one.

If the Event ID exceeds 2.1 billion, the event service will be unable to save new events. If this occurs:

1. **Immediately** open the Windows Services Manager (services.msc) and locate the **Protege GX Event Service**.
2. Right click and select **Properties**. Set the **Startup type** to *Disabled*.
3. Click **Apply**.
4. Click **Stop**.
5. Repeat for **Protege GX Extended Services Events** to stop the extended event service as well (if installed).

While the event services are stopped, controllers will save incoming events to prevent them from being lost (up to 50,000 events per controller).

6. Contact ICT Technical Support as soon as possible for assistance with backing up your events database and creating a new one.

Integrations

Enterprise-level systems typically have one or more integrations for video cameras, elevators, user management, reporting and other functions. Some sites also create unique user interfaces such as web clients and mobile apps that access Protege GX data.

When using integrations created by ICT or a third-party software provider:

- Engage with both ICT and the third-party provider to understand your licensing obligations for both systems.
- Read the application note or other documentation for the integration before starting deployment.
- Ensure that you meet the system requirements for both Protege GX and the third-party software. Typically you should run the two servers on different computers to ensure that you have capacity.
- Ensure that there is a stable network connection between the two servers. Consult with the site's IT team to set up the required networking routes and rules before you begin.

When creating custom integrations for Protege GX:

- Always connect to Protege GX using a validated method such as the Protege GX SOAP Service.
Contact ICT for access to the *Protege GX SOAP Service API Specification*.
- Apply rate limiting and other best practices for API usage.
- Never directly modify the Protege GX databases. This method bypasses the Protege GX services, which can cause unexpected issues with the system. ICT cannot provide support for issues caused by modifying the database directly.
- Always validate integrations on the test bench before deploying them to live sites.

Programming

This section contains some recommendations for programming and managing large systems.

Naming Conventions

As the site grows, naming conventions become increasingly important to prevent confusion. Ensure that there is a naming convention in place **before** starting to program the site. A good convention uses common elements such as location to help operators sort and filter records.

For example, a door for a server room in the Auckland office might be called: AKL - FL02 - Server Room Door. This format has a number of benefits:

- Abbreviations keep the name brief, making it easier to read in status pages and events.
- When operators sort the doors page or status list, doors in the same physical location will be grouped together.
- An operator can find all doors in the Auckland office by searching for "AKL".
- An operator can find all doors on Floor 2 by searching for "AKL - FL02".

Pagination and Searching

Large systems typically have very many users, access levels and doors. This can increase the loading times of these pages. We recommend that operators keep the default pagination setting of 200 records per page to ensure that loading times are reasonable.

For these pages, the quick search bar above the record list is the quickest way to search for records by name. The Find tool is also available for searching based on other settings.

Access Levels and Groups

When an access level needs over 100 doors, we recommend using door groups to assign these to the access level rather than assigning individual doors. The server and controllers can handle large door groups more efficiently than access levels with large numbers of individual doors.

Floor Plans

Floor plans are a useful tool for viewing and controlling large numbers of devices.

A good way to manage a large system is to break the site into multiple floor plans (e.g. one per building floor), then create a 'directory' or 'site map' floor plan. The directory contains buttons that link to each floor plan, making it quick to 'travel' around the site.

For example, for a large university campus you might use three levels of floor plans:

1. One high-level site map that shows the buildings on the campus. This floor plan does not include any devices, but has a button for each building.
2. One plan for each building, showing all the floors. These floor plans do not include any devices, but have a button for each floor.
3. One floor plan for each floor of the building. These floor plans include all of the devices for their floor. They can also include buttons that link to adjacent floors, and back to the building plan.

Designers & manufacturers of integrated electronic access control, security and automation products.

Designed & manufactured by Integrated Control Technology Ltd.

Copyright © Integrated Control Technology Limited 2003-2026. All rights reserved.

Disclaimer: Whilst every effort has been made to ensure accuracy in the representation of this product, neither Integrated Control Technology Ltd nor its employees shall be liable under any circumstances to any party in respect of decisions or actions they may make as a result of using this information. In accordance with the ICT policy of enhanced development, design and specifications are subject to change without notice.