



Pythia Technologies

DM Console

User Guide

DM · SPDM · SPIGM · SPAC4

Designed for AI-assisted answers — upload this document to your AI assistant and ask.

Last updated 2026-07-28

How to Use This Document

This guide is designed to be used with an AI assistant. You do not have to read it from front to back. Instead, give this document to an AI chat assistant and simply ask it your question in plain language — it will answer using the information in these pages, walk you through the steps, and explain anything you do not understand.

Most people never read a manual cover to cover, and they should not have to. Everything about the Pythia **DM family** — the **DM**, **SPDM**, **SPIGM**, and **SPAC4** products — is written down here in full detail so that an AI assistant can read it for you and give you a direct answer the moment you need one. It is one guide for all of the products; the assistant will pull the right answer for whichever one you have.

Quick start

1. **Save this document** to your computer, phone, or tablet.
2. **Open an AI assistant** you have access to. Any of the common ones work — for example Claude (claude.ai), ChatGPT (chatgpt.com), Microsoft Copilot, or Google Gemini.
3. **Start a new conversation and attach this document.** Look for a paperclip, +, or **Attach** button and select this file. If your assistant cannot accept file attachments, open the document, copy all of its text, and paste it into the chat instead.
4. **Ask your question in your own words**, then ask follow-up questions as they come up — just like talking to a knowledgeable coworker.

Note: If your AI assistant offers a choice of file, the plain-text (.md) version of this document uploads and reads most reliably. A PDF or Word version works too; the words are the same in all of them.

Questions you can ask

You can ask anything covered in this guide. A few examples to get you started:

- "How do I add a new location, group, and unit?"
- "What does the red **COMMUNICATIONS** banner on a unit mean, and what should I check?"
- "Set me up to email the on-call team, but only for critical alarms."
- "Walk me through changing the gateway's IP address safely."
- "A backup will not download — what do I check?"
- "How do I read the trend chart, and how do I turn on trending for a data point?"
- "What is the difference between an alarm setting and an alarm event?"
- "How do I give a coworker a View-only login?"
- "On my SPDM, how do I bring a Smart Panel input into a unit?"

- "On SPAC4, how do staging and failover decide which units run?"
- "Explain this message I am seeing: *(paste the message)*."

Tips for good answers

- **Be specific.** Name the page, unit, or message you are working with. "The Add New Unit dialog rejected my driver" gets a better answer than "it will not save."
- **Ask for step-by-step** when you want to follow along: "Give me the exact steps."
- **Paste the message you see on screen** and ask what it means and what to do about it.
- **Ask follow-ups.** If an answer is not quite what you needed, say so and the assistant will refine it.
- **Ask where to find something:** "Where in the console do I configure heartbeat emails?"
- **Say which product you have** — DM, SPDM/SPIGM, or SPAC4 — if you are not on a plain DM. It helps the assistant give you the right product-specific steps instead of the general ones.

What the assistant can help with

Signing in and moving around the console; monitoring your sites, groups, and units; understanding equipment dashboards; reading and clearing alarms; setting up email and SNMP-trap notifications and heartbeats; using trends; managing users and passwords; system settings such as time, backups, software updates, logs, and network configuration; and connecting other systems over SNMP, Modbus, and BACnet. It also covers the product-specific features — the Smart Panel and DIRS input-to-unit mapping on **SPDM/SPIGM**, and the staging, failover, and auto-sequencing controls on **SPAC4**.

> **Important:** Do **not** paste passwords, real customer or site names, IP addresses, or other sensitive details into a public AI tool unless your organization has approved that tool for such use. This document itself contains no site-specific secrets, and you do not need to share any to get help with how the console works. The assistant answers from this document; if your question is not covered here, contact Pythia Technologies support rather than acting on a guessed answer.

You can still read it yourself

If you prefer to read, the full guide follows. A table of contents lists every chapter, and each chapter is organized by task. Whether you read it or let an AI assistant read it for you, the answers are the same.

How This Guide Is Organized

This guide covers the Pythia **DM family** of products — **DM**, **SPDM** (and **SPIGM**), and **SPAC4**. They are built on one platform, so most of what you need to know is the same across all of them. Rather than repeat that shared material three times, the guide is arranged in two layers.

- **Part 1 — Core Platform** describes everything the products have in common: signing in and moving around the console, monitoring your locations, groups, and units, equipment dashboards, alarms, automations, trends, email delivery, users, system settings, and integrations. **Part 1 applies to every product.**
- **Parts 2–4** are short, product-specific sections that cover only what is unique to **DM**, **SPDM/SPIGM**, and **SPAC4**.

How to find your answer

Start with the product you are using. If that product's section covers your topic, use it. For anything the product section does not mention, the **Core Platform** (Part 1) applies. A product section only ever **adds to** or **overrides** the core behavior — it never repeats it.

There is one important exception. **SPAC4** is a controller, not a general-purpose monitor: it has no Locations tree and uses a different editing model, so it **replaces** several core concepts rather than adding to them. The SPAC4 section states plainly what from the Core Platform does **not** apply to it.

> **Note for AI assistants using this document:** Treat **Part 1 (Core Platform)** as the default source of truth for every product. When a question names a specific product, first check that product's section (Parts 2–4) for anything unique or overridden, then fall back to Part 1 for everything else. If a reader asks about a "SPAC4 feature" or an "SPDM feature" that is really shared behavior, answer from Part 1 and note that it applies to their product too.

Which product do I have?

- **DM (Data Monitoring)** — the full monitoring platform. You build a tree of **Locations** → **Groups** → **Units** and monitor equipment over SNMP, Modbus, BACnet, and other drivers. Part 1 **is** the DM.
- **SPDM / SPIGM (Smart Panel Data Monitor)** — a DM built around a wired **Smart Panel** (a CM06P controller). The panel's inputs become ordinary Units through the **DIRS** driver; everything else is the Core Platform.
- **SPAC4 (Smart Panel AC-4 Controller)** — a **Smart Panel** that **controls** air-handling equipment: a fixed panel that runs staging, failover, and auto-sequencing for up to eight units. It drives equipment rather than only watching it.

If you are not sure which product you have, the sign-in banner and the shape of your navigation tree are the quickest tells: a Locations tree means DM or SPDM; a fixed **Smart Panel** → **Group 1 / Group 2** layout with no Locations means SPAC4.

Part 1 — Core Platform

This section applies to every DM-family product (DM, SPDM, SPIGM, and SPAC4). Anything unique to a specific product is in Parts 2-4; where a product section is silent, the behavior described here applies.

Chapter 1 — Introduction

Welcome to the Pythia Technologies DM Console User Guide. The DM is a Windows-based monitoring and control gateway that watches over your power and environmental equipment — UPS systems, generators, battery systems, air conditioning units, temperature and humidity sensors, power panels and more — and presents everything through a web console you open in an ordinary browser. This chapter introduces the ideas the rest of the guide builds on: how your equipment is organized, how alarms work, and what the two user roles can do. There are no procedures here; every concept introduced below gets its own chapter later in the guide.

What the DM is

The DM sits on your network between your equipment and you. It continuously polls each monitored device over the device's own protocol, evaluates the readings against the alarm settings you configure, records history, and sends notifications when something needs attention. Where a device supports it, the DM can also send commands back — switching an output on or off, or adjusting a setpoint.

You interact with all of this through the DM console: a web page served by the gateway itself. There is nothing to install on your workstation — you sign in from a browser, and everything you can see or configure is on screen.

The console has four fixed regions:

- The **navigation tree** on the left mirrors your site's structure. Click any item to open its page.
- The **top bar** holds global search, **Help**, and **Sign Out**.
- The **content area** in the middle shows the page for whatever you selected.
- The **status bar** along the bottom shows where you are, a site-wide alarm badge, the system clock, and who is signed in.

Chapter 2 walks through the shell in detail; the rest of this chapter is about what fills it.

How your site is organized: Locations, Groups, Units

Everything the DM monitors lives in a three-level hierarchy:

Level	What it represents	Example
Location	A physical site or building	Sunbury Office
Group	An area or category within a location	Environmentals, Power Units
Unit	A single monitored device	Office TH1, CRAC-01, UPS-21A

A location contains groups, and a group contains units. The navigation tree reflects this exactly: expand **Sunbury Office** and you see its groups; expand **Power Units** and you see units such as **UPS-21A**,

Genset 1, and **Battery System**. How you divide things up is entirely yours — groups can represent rooms, equipment categories, or anything else that makes your site easy to navigate.

Administrators create, rename, and delete locations, groups, and units; Chapter 4 covers that. For everyone else, the hierarchy is simply the map you browse.

What a unit contains

A unit is where monitoring actually happens. Each unit carries three kinds of things:

- **Data points** — the live readings the DM collects from the device: a temperature, a battery voltage, an input status, an output load percentage. Data points update automatically as the console refreshes, and each one can be charted and exported.
- **Controls** — commands the DM can send to the device, where the device supports them: on/off outputs and numeric setpoints. Not every unit has controls; a simple sensor is read-only.
- **Alarm settings** — the rules that decide when a reading is a problem: which data point to watch, what threshold or condition triggers, and how severe the resulting alarm is.

Chapter 5 covers all three in depth. Many common device types also get a dedicated dashboard page — a live equipment diagram rather than a plain list — which Chapter 6 tours.

Alarm settings vs. alarm events

Two related terms are easy to conflate, and the distinction matters throughout this guide:

- **Alarm settings** are **configuration**, and they live on each unit. They answer the question "what should count as a problem for this device?" — for example, "alarm when **Office TH1** temperature exceeds 80". Administrators define them; they sit quietly until a reading crosses the line.
- **Alarm events** are **occurrences**, and they live in the site-wide alarm buckets. When a reading violates an alarm setting, the DM raises an event with a timestamp, a severity, and a message telling you what happened and where.

Alarm events are collected into three site-wide buckets:

- **Active** — alarm conditions happening right now. This is the list to work when something is wrong.
- **Information** — informational events that don't represent an ongoing fault.
- **History** — the permanent record of past events, searchable and exportable.

Each event carries a severity level (Critical, Major, Minor, Warning), each with its own color, so you can tell at a glance how urgent an entry is. Chapter 7 covers the alarm pages, severities, and how to work the lists; the alarm settings that generate the events are configured per unit in Chapter 5.

Automations: getting notified

Watching the console is one way to catch an alarm; having the DM tell you is better. **Automations** are the DM's notification rules. An automation can send email to one or more recipients and can send

SNMP traps to your network management system when alarm events occur. Automations also support **heartbeat** messages — scheduled proof-of-life emails that confirm the DM itself is up and reachable, so silence never gets mistaken for "all normal". Time exceptions let you suppress notifications during defined windows, such as scheduled maintenance. Chapter 8 covers automations end to end.

Trends: recorded history

Live readings tell you what is happening now; **trends** tell you what has been happening. The DM can record selected data points into trend history and chart them over time — useful for spotting a slowly warming room, tracking UPS load growth, or documenting conditions after an incident. You choose which data points to record, then read, zoom, and export the resulting charts. Chapter 9 covers trending.

Status rollup: how the console shows health

Every unit, group, and location has a status, and status flows **upward**:

- A **unit** is **Normal** when all its readings are within limits, and **Alarm** when any of its alarm settings has triggered.
- A **group** shows **Alarm** whenever any unit inside it is in alarm.
- A **location** shows **Alarm** whenever any of its groups is in alarm.
- The **navigation tree** colors every level to match — green icons for Normal, red icons and red labels for Alarm, gray for Disabled.

A unit, group, or location can also be **Disabled** — deliberately excluded from monitoring, for instance during maintenance. Disabled items show gray and are never counted as alarms, even if their last readings were out of range.

The practical effect: one sensor over its limit in **Environmentals** turns that unit red, turns **Environmentals** red, turns **Sunbury Office** red, and lights the red **N active alarms** badge in the status bar. You can spot trouble from anywhere in the console and follow the red down the tree to the exact device. When nothing is wrong, the status bar shows a green **All normal** badge instead.

The two user roles

Every account on the DM has one of two access levels:

- **Administrator** — full configuration rights: add and delete locations, groups, and units; define alarm settings; send control commands; manage automations, trends, users, and system settings.
- **View** — read-only monitoring: browse every page, watch live data, search, open property dialogs to read settings, and export reports — but no buttons that change anything.

The console doesn't gray out administrator controls for View users; it hides them entirely. If you sign in with View access and a button described in this guide isn't on your screen, that is why. Sections of this guide that require Administrator access are marked with an **Admin only** callout, and each chapter notes

what View users see instead. Your current role shows next to your username in the status bar. Chapter 11 covers accounts and user management.

Note: Control commands are Administrator-only for good reason — they act on real equipment. Chapter 5 explains the safeguards (and the lack of a confirmation step) in detail.

How this guide is organized

- **Chapter 2 — Getting Started:** signing in, the console shell, navigating with the tree, global search, and signing out.
- **Chapter 3 — Working with the Console:** the controls shared by every page — Cards and Table views, auto-refresh, property dialogs, add and delete flows, and connection indicators.
- **Chapter 4 — Monitoring Your Site:** the site overview and the Locations, Groups, and Units pages, plus creating and managing the hierarchy.
- **Chapter 5 — Units in Detail:** the unit page, live data and charts, controls and setpoints, and per-unit alarm settings.
- **Chapter 6 — Equipment Dashboards:** the dedicated live pages for UPS, generator, battery, CRAC, power distribution, and sensor equipment.
- **Chapter 7 — Alarms:** the Active, Information, and History buckets, severities, and the wall-display alarm view.
- **Chapter 8 — Automations & Notifications:** email and SNMP trap notifications, time exceptions, and heartbeat proof-of-life messages.
- **Chapter 9 — Trends:** assigning data points to trend history and reading and exporting trend charts.
- **Chapter 10 — Smart Panel (SPDM):** monitoring and configuring the Smart Panel input module.
- **Chapter 11 — Users & Accounts:** your own account settings and, for administrators, managing users and reviewing the logon audit log.
- **Chapter 12 — System Settings:** license, date and time, backups, software updates, system logs, password policies, and network configuration.
- **Chapter 13 — Integrations (SNMP · Modbus · BACnet):** exposing DM data to building management and network management systems through protocol device mappings.

Three appendices follow: a reference of user-visible messages (Appendix A), an access-rights matrix comparing View and Administrator page by page (Appendix B), and a glossary of DM terms (Appendix C).

Chapter 2 — Getting Started

This chapter gets you from a browser address to a working console session. You will sign in, learn your way around the four areas of the console window — the navigation tree, the top bar, the status bar, and the content area — run your first search, and sign back out. Everything in later chapters assumes the layout and navigation habits introduced here.

Opening the console

The DM console runs entirely in a web browser — there is nothing to install on your computer. Your administrator will give you the address of your DM, for example:

```
`` http://SITECODE/ ``
```

Type the address into any modern browser. If you are not already signed in, the console takes you straight to the sign-in page.

Note: JavaScript must be enabled in your browser. If it is off, the sign-in page shows a red notice — "JavaScript is required to sign in to this system." — instead of the form. Enable JavaScript and reload the page.

Signing in

The sign-in page is a single dark card with the Pythia Technologies banner, a **User** field, a **Password** field, and a **Sign On** button.

To sign in:

1. In the **User** field, type your sign-in email address (for example `user@example.com`). Capitalization does not matter.
2. Type your password in the **Password** field. To check what you typed, click **Show** next to the field — the password becomes visible and the button changes to **Hide**. Click **Hide** to conceal it again.
3. Click **Sign On**, or press **Enter** from either field.
4. The button reads **Signing in...** for a second or so, and then the console opens.

That brief **Signing in...** pause is normal and is actually a security feature: your browser spends that moment converting your password into a one-time cryptographic proof, so the password itself is never sent across the network in readable form.

Note: If amber text reading "Caps Lock is on." appears beneath the password field, your Caps Lock key is active — passwords are case-sensitive, so turn it off before continuing. The warning disappears on its own when Caps Lock goes off.

If sign-in fails

If the sign-in is not accepted, a red message box appears and the password field is cleared automatically so you can re-enter it.

- **"Invalid credentials"** — the user name or password was not accepted. This message is deliberately generic: for security, the console never reveals whether the account exists, the password was wrong, or the account has been disabled. Check your typing (and Caps Lock) and try again. If you still cannot get in, contact your administrator.
- **"Nonce expired"** — you left the page sitting for more than a minute between loading it and clicking **Sign On**. Simply try again; the retry succeeds immediately.
- "Enter your email or username." / "Enter your password." — you left a field blank.

There is no limit on sign-in attempts, so you can keep trying without fear of locking the account.

The one-time email update prompt

Sign-in accounts on the DM are email addresses. If your account was created on an earlier version of the software with a short user name (such as `operator`), the first time you sign in after an upgrade the console stops at an **Update Your Login** page instead of opening. An amber banner explains: "Please update your account to use an email address as your login. This is required before you can continue."

The page shows your **Current login** (read-only) and asks for a **New email**. To complete the update:

1. Type your email address in the **New email** field.
2. Click **Change Email**. The button reads **Saving...** while the account is updated.
3. A green **Email Updated** panel confirms the change, shows your new address, and tells you: "You have been signed out. Click OK to sign in again using your new email."
4. Click **OK, Sign In** to return to the sign-in page, then sign in with your **new email address and your existing password** — your password has not changed.

If the address you enter is not accepted — for example "That email is already in use by another account" — the page explains why and lets you try again.

Note: Clicking **Cancel** (or the **X**) does not skip this step — it signs you out and returns you to the sign-in page, and your next sign-in with the old name lands right back on this prompt. You cannot use the console until an email login is set.

This happens once per account. After the update, you sign in with the email address from then on.

The console layout

After sign-in, the console window is divided into four permanent areas:

Area	Where	What it does
Navigation tree	Left column	Every location, group, unit, and settings page in your system, arranged as an expandable tree. The primary way to move around.
Top bar	Across the top	The global search box, the Help button, and the Sign Out button.
Content area	The large center-right region	The page for whatever you selected in the tree.
Status bar	Across the bottom	Your current position (breadcrumb trail), the console-wide alarm badge, the system clock, and your user name with a role badge.

The tree, top bar, and status bar are always present; only the content area changes as you navigate.

The navigation tree

The tree's top row shows the product logo and your site's name — clicking it takes you to the site overview. Below that, every row is one node in your system: an icon, a name, and — when the node contains children — a chevron (>) to its left.

Two different clicks do two different things:

- **Click a node's name** — the node is selected (highlighted, with a blue edge on its left) and its page loads in the content area. If the node has children, its branch also expands.
- **Click the chevron** — the branch expands or collapses **without changing the page**. Use chevrons to look around the tree while staying on your current page.

Which branches you have open is remembered as you navigate, so the tree stays the way you arranged it for the rest of your browser session.

The tree also tells you system health at a glance:

- A **red** name and icon mean there is an active alarm on that node **or somewhere below it** — a red location contains at least one unit in alarm. Follow the red down the branches to find the source.
- A **gray, dimmed** name means the node is disabled and is not being monitored.
- Normal nodes show a green icon.

These colors update automatically while you work — you do not need to reload anything to see a new alarm appear in the tree.

The top bar

From left to right, the top bar holds:

- The **global search box** — see "Finding things with search" below.
- The **Help** button (question-mark icon) — opens this guide in a new browser tab.
- The **Sign Out** button — ends your session (see "Signing out" below).

Occasionally a white pill with a red border appears here reading **Update available — Refresh**. It means a newer console version was installed on the DM while you were signed in. Click the pill: the console reloads all of its areas and you continue where you were, now on the new version. Nothing is lost — this is just a page refresh, not a sign-out.

The top bar is identical for every user; nothing in it depends on your access level.

The status bar

The bottom strip answers three questions at all times: *where am I, is anything wrong, and who am I signed in as.*

Left — the breadcrumb trail. After the label **DIRECTORY:**, the bar spells out the path to the page you are viewing, for example / Sunbury Office / Environmentalals / Office TH1. Every earlier segment is a blue link — click one to jump straight to that level (the leading / returns to the site overview). The last segment is the page you are on, marked with a status dot: green for **Normal**, red for **Alarm**, gray for **Disabled**.

Right — the status cluster:

- **The alarm badge.** With no active alarms it is a quiet green pill reading **All normal**. When alarms are active it turns red and shows the count — for example **3 active alarms** — and becomes a shortcut: click it to jump directly to the Active alarms list (Chapter 7). If the badge instead reads **Disconnected – retrying**, your browser has temporarily lost contact with the DM; it reconnects on its own, so no action is needed unless it persists.
- **The clock.** The DM's system time, ticking live, with the date beside it. Alarm and event timestamps throughout the console use this clock.
- **Your identity.** Your sign-in name followed by a role badge: **Admin** in blue, or **View** in gray. If you ever wonder why a page shows no editing buttons, glance here first — View users see the console read-only.

The content area

Everything else on screen is the content area: dashboards, live data, settings pages. Its contents always correspond to the selected tree node and the breadcrumb trail. When you follow a link *inside* a content page — clicking a unit card on a group page, say — the tree selection and breadcrumb move along with you, so the three always agree. Chapter 3 covers the controls that content pages share (view toggles, auto-refresh, the Properties modal, and more).

Finding things with search

When you know the name of what you want — a unit, a location, an alarm — search is faster than walking the tree.

1. Click in the search box in the top bar (placeholder text: "Search alarms, locations, groups, units...").

2. Type at least **2 characters** of the name.
3. Press **Enter**. The search results page opens in the content area.

Results are grouped by type — **Units, Locations, Groups, Alarms, Controls, Data Points** — with a summary line showing the match count. Each result row carries a status dot (green, red, or gray, just like the tree) and, where useful, a muted line of context showing where the item lives. Click any result to go to it; the console navigates the tree and opens the item's page.

On the results page itself:

- The search box at the top is pre-filled with your query, and the results **refine live as you type** — no need to press Enter again.
- Pressing **Enter** in that box jumps straight to the first result.
- A very broad query may be capped, with the summary reading "Showing the first N of TOTAL matches — refine your search." Add another character or two to narrow it down.
- If nothing matches you will see, for example, "No matches for "zzzz"."

Getting help

Click **Help** in the top bar at any time. This guide opens as a PDF in a new browser tab, leaving your console session untouched.

Signing out

Click **Sign Out** at the right end of the top bar. Your session ends immediately and the browser returns to the sign-in page — there is no confirmation step. Sign out whenever you leave a shared workstation; anyone at the browser while you are signed in has your level of access.

You are now signed in, oriented, and able to find your way around. Chapter 3 introduces the controls that recur on almost every console page, and Chapter 4 begins the tour of your site's monitoring pages.

Chapter 3 — Working with the Console

Every page in the DM console is built from the same small set of controls: a Cards/Table view switch, an auto-refreshing data area with a status dot, a **Properties** dialog for editing, an **Add** dialog for creating, and a confirmation dialog for deleting. This chapter teaches those patterns once, so the rest of the guide can say "open **Properties**" or "click the rename pencil" without re-explaining what happens. If you are comfortable with the material here, you already know how to operate most of the console — later chapters mainly tell you *which* fields each page offers.

Cards and Table views

Most list pages — locations, groups, units, settings areas, and any other container — can show their contents two ways:

- **Cards view** — a grid of clickable tiles. Each card shows the item's icon and name in bold, and, when the item has a description, a second muted line with a small info icon. Hovering a card gives it a blue border; clicking anywhere on the card opens that item.
- **Table view** — one row per item. The columns are specific to the page (a units list shows different columns than a settings list), with the item's icon in the first column. Clicking a row opens that item.

Switch between them with the two buttons at the right end of the section header: **Cards** (grid icon) and **Table** (lines icon). The active view renders as a solid blue button.

Note: Your choice is remembered on this browser and applies to **all** container pages, across sign-ins. Pick **Table** once and every list page opens in Table view until you switch back.

Both views behave identically otherwise — same items, same navigation, same toolbar. If a page has nothing in it, both views show "No items in this container."

Auto-refresh, the status dot and the Updated stamp

Content pages keep themselves current — you never need to press the browser's reload button. Each page re-fetches its data automatically on a repeating cycle:

- **Fast pages** — live equipment dashboards and other real-time views refresh every **2 seconds**, so readings track the device closely.
- **List pages** — container lists, alarm lists and similar pages refresh every **30 seconds**.
- The list-page rate is configurable system-wide by an administrator; pages that use it show a footer note such as "Auto-refresh every 30 seconds". Setting the rate to 0 turns polling off ("No auto-refresh") — pages then load once and stay static until you navigate.

Two indicators at the top-right of every content page tell you the refresh machinery is healthy:

- **The status dot** — green and pulsing when the page is healthy, amber while data is loading, steady red when the last fetch returned an error, gray when the item is disabled.
- **"Updated: HH:MM:SS"** — the time of the last successful data fetch, stamped next to the dot. If this time keeps advancing, you are looking at live data.

A few pages are deliberately static — for example, the leaf properties pages described below load once and note "Property values are static on this page."

The Properties dialog

Properties is the console's single editing pattern. Anything you can configure — a location, a unit, a data point, an email recipient, a system setting — is edited through a Properties dialog with the same layout everywhere.

Admin only: The **Properties** button (and everything in this section) appears only for Administrator accounts. View users see the same pages without any editing controls.

You reach it two ways:

1. **From the toolbar** — container pages show a blue **Properties** button in the section header. It edits the item whose page you are on.
2. **From a row or link** — on some pages, clicking an item's name opens the Properties dialog for that item directly, so you can edit a child without leaving the list.

Some simple items — a single email recipient, an allowed IP address — get a whole page of their own: a three-column **Property | Value | Description** table with a **Properties** button above it. The dialog it opens is the same one you meet everywhere else.

Field types you will meet

The dialog shows one field per property. The control depends on the kind of value:

You see	It means
Text or number box	Type a value directly. Number boxes accept whole numbers or decimals as appropriate.
Toggle switch	A Yes/No setting. Green and to the right = Yes (on); gray and to the left = No (off). Read-only tables display these values as Yes / No .
Dropdown	Pick one value from a list. Some dropdowns include a blank first entry — choosing it clears the value.
Checkbox list	A scrollable multi-select — tick every option that applies (for example, choosing several alarm categories at once).
Password box	Always shown empty with the placeholder "Leave blank to keep current". The stored password is never displayed. Leaving it blank keeps the existing password; typing a value replaces it.

Editable fields have a slightly brighter background than read-only ones. **Read-only fields** (live readings, system-assigned values) appear grayed and cannot be changed; a read-only field with no value at all is simply hidden. In rare cases you may see a disabled field with a yellow note "Property type N not yet editable in this dialog" — that value can be viewed but not yet edited from this dialog.

Saving

Click **Save** to write all your changes in one step. The button disables briefly while the save runs (a double-click guard); on success the dialog closes and the page refreshes to show the new values. If the server rejects the save, its message appears in red at the bottom of the dialog and nothing is changed — correct the value and try again.

One rule applies to every field in every dialog: values may not contain the pipe character (|). Trying to save one is blocked with a message naming the field, for example: Field "Description" cannot contain a pipe character (|).

Renaming with the pencil

Where an item's name can be changed, the Properties dialog shows the name at the top with a small **pencil** button beside it. Renaming is deliberately separate from **Save** — it takes effect on its own.

1. Click the pencil. The name becomes an edit box with a green **checkmark** (confirm) and a gray **X** (cancel) beside it. Enter also confirms; Esc cancels.
2. Type the new name. The checkmark lights up once the value differs from the current name.
3. Click the checkmark. The name changes immediately — in the dialog, in the navigation tree and in the page address. Everything beneath the renamed item keeps working.

If the new name is not acceptable, a red message appears inline next to the box and the name is not changed:

- Names must be 1–100 characters: "Name is required" / "Name must be 1-100 characters".
- Names may not contain any of these characters: \ & % # ? " , — the message reads "Invalid characters: \ & % # ? " , ,".
- A sibling with the same name already exists: "That name already exists".

Some items are **named after** the thing they represent, and the console checks the format too:

- An item named by an **IP address** must be a valid dotted-quad: "Must be a valid IPv4 address (e.g. 192.168.1.10)".
- An item named by an **email address** must look like one: "Must be a valid email (e.g. user@example.com)".

Items whose names must not change show the name as a grayed, disabled box with no pencil.

Add dialogs

Container pages that accept new items show a blue **Add** button in the toolbar. It opens a dialog titled **Add to** followed by the container's name, with one input per field — the same field types as the Properties dialog.

- ****Required fields are marked with an asterisk (*)**** — the dialog will not save until they are filled in.
- Validation messages appear in red **under the offending field** before anything is sent to the server. For example, entering a malformed IP address produces: "IP Address" must be a valid IPv4 address (e.g. 10.251.0.0) -- check for an extra dot or space.
- The pipe character is rejected here too: Field "Name" cannot contain a pipe character (|).

Click **Save** to create the item. On success the dialog closes, a brief **"Saving..."** overlay appears (see below), and the console navigates straight to the newly created item so you can carry on configuring it.

Deleting items

Delete is never a toolbar button. It always lives **inside the Properties dialog**, as a red **Delete** button in the bottom-left corner — a small speed bump that keeps destructive actions one deliberate step away from the list.

Clicking **Delete** closes the Properties dialog and opens a confirmation dialog with a red title, naming exactly what will be removed, with the caution "This action cannot be undone." Click the red **Delete** button to proceed, or **Cancel** to back out.

Caution: Deletion is permanent. There is no undo and no recycle bin. When the confirmation includes the phrase "and everything inside it", the delete is a **cascade**: the item **and all of its contents** — child containers, their settings, their history — are removed together.

Some containers refuse to cascade and instead require you to empty them first. In that case the confirmation is replaced by a refusal explaining what is still inside. For example, attempting to delete a location that still has groups shows:

- > The location "Sunbury Office" contains 2 groups.
- > Delete all groups before deleting this location.

Delete the children first (or move your intent elsewhere), then delete the container. Generic containers word the same refusal as "Container holds N child container(s). Remove children before deleting."

On a successful delete, the console shows the **"Saving..."** overlay briefly and then takes you to the deleted item's parent page.

The "Saving..." overlay

After you add, rename or delete something, the whole screen briefly dims under a dark overlay with a spinner and the label **"Saving..."**. This is normal: the server needs a moment to finish applying the change before the new (or renamed, or surviving parent) page can be displayed, and the console is simply waiting for it to be ready.

There is nothing to click and nothing to worry about — the overlay clears itself, typically in a second or two, and the console then navigates to the right place on its own. Just wait it out.

When someone else deletes what you're viewing

More than one person can use the console at once. If another user deletes the item you currently have on screen, your page's next auto-refresh detects it and replaces the content with a notice and a live countdown:

- > Redirecting to *parent* in 3 second(s)... **Go now**

Auto-refresh stops, the countdown ticks down from 3, and when it reaches zero the console navigates you to the deleted item's parent page, which still exists. Click **Go now** if you would rather not wait.

Nothing on your side caused this, and nothing needs fixing — it is simply the console keeping every signed-in user's view consistent.

Connection indicators

Two indicators tell you when the console cannot reach the DM — during a restart, a network interruption, or a software update:

- **The "Connection lost — reconnecting..." pill.** When a page's refresh fails at the network level, an amber-bordered pill with a spinner floats at the bottom-center of the content area. The page keeps retrying by itself; the pill disappears on the first successful response. If a page is **first opened** during an outage, it shows an inline panel instead: a spinner with "Connection lost — attempting to reconnect..." and "This page will refresh on its own once the connection is back."

- **The status bar "Disconnected – retrying" badge.** If the status bar cannot reach the server for roughly 12 seconds (long enough to ride out one or two dropped polls), its alarm badge turns red and reads **Disconnected – retrying**, with the tooltip "Cannot reach the server". The breadcrumb's status dot also turns red.

Both indicators recover on their own as soon as the server answers again — no reload needed. If they persist, treat it as a real outage: check the DM's power and network connection.

The "Update available" pill

When an administrator installs a new console version on the DM, a pill appears in the top bar: white with a red border, a pulsing red refresh icon, and the text **Update available — Refresh**. Your current pages keep working, but they are running the previous version's code.

Click the pill when convenient. The console reloads all of its frames — content, tree, top bar and status bar — so you pick up the new version, and the notice is acknowledged for your account on this browser. Other users see their own pill until they refresh too.

Export buttons

List-style pages that hold reportable data — alarm lists, data readings, mapping tables and others — carry export buttons, typically **PDF** and **Excel** (some pages offer more formats). Exports generate entirely on the DM and download straight to your browser; no internet connection is required. What each export contains is specific to its page, so the chapter covering that page describes it.

What View users see

Everything in this chapter that *shows* information works identically for View users: Cards and Table views (including the remembered preference), auto-refresh, the status dot, property tables, connection indicators, the update pill, and exports. What View users never see are the editing controls: no **Add** or **Properties** toolbar buttons, no rename pencil, no **Delete** — the buttons are absent entirely, not merely disabled. Where an Admin's page footer reads "Use Properties to edit or delete this item.", a View user's reads "Property values shown in read-only mode."

Chapter 4 — Monitoring Your Site

The DM organizes everything it monitors into a three-level tree: **Locations** (physical sites or buildings), **Groups** (areas or categories within a site), and **Units** (individual monitored devices). This chapter walks you through the pages that present that tree — from the Site Overview landing page down to the units in a single group — and shows administrators how to build and maintain the hierarchy: adding, renaming, and deleting locations, groups, and units. Operators with View access can browse everything described here; the management procedures near the end of the chapter are Admin-only.

The pages in this chapter share the common controls described in **Working with the Console** — the **Cards/Table** view toggle, auto-refresh and the status dot, the Properties modal pattern, and the "Saving..." overlay. Those are not repeated here; where a page departs from the common pattern, the difference is called out.

How status rolls up

Every location, group, and unit carries one of three status states, shown as a colored badge on its card or row:

- **Normal** (green) — the item is enabled and reporting no alarms.
- **Alarm** (red) — the item, or any unit beneath it, is in alarm.
- **Disabled** (gray) — the item has been turned off for monitoring. A disabled item is never counted as an alarm, even if its last known state was Alarm — disabled always wins.

Status rolls **up** the tree: when any unit goes into alarm, its group and its location go into alarm too. That means a red badge on the Site Overview page always leads somewhere — click through the red items and the tree takes you straight to the unit that raised it.

Note: If a device reports a status the DM does not recognize, the badge shows that text in gray (for example **Unknown**) and the item is counted with the Normal total, not the Alarms total.

The Site Overview page

The Site Overview is the DM's landing page — the page you see when you sign in or click the topmost node in the navigation tree. It is designed as a springboard: identify your site, spot any alarms at a glance, and click down into a location.

The top half of the page shows your site image — a photograph or diagram supplied for your installation — with a floating label reading **SELECT A LOCATION BELOW**. The bottom half lists your locations:

- **Card view** (the default) shows one card per location: an icon and the location name, an optional description, and a status pill on the right — green **Normal**, a red pulsing **ALARM**, or gray **Disabled**.

Cards in alarm get a red border and a reddish background so they stand out across the room.

- **Table view** shows the same locations in three columns: **Location** | **Status** | **Description** (the description column shows -- when empty). Rows in alarm are tinted red.

Click any location card or row and the navigation tree moves to that location, which opens its Groups page.

A few behaviors are specific to this page:

- It refreshes automatically every 30 seconds, regardless of the site's configured refresh rate.
- The **Cards/Table** choice is not remembered here — the page always opens in card view. (Every other page in this chapter remembers your choice; see *Working with the Console*.)
- The page is deliberately forgiving about brief network hiccups: once the location list has been drawn, a failed refresh never blanks it. Only when the page cannot load anything at all — two failures in a row on a fresh visit — does it show the banner "⚠ Error Loading Data" with the text "Unable to load location data. Check the API connection."

Admin only: A **Properties** button in the Locations section header opens the **Site Properties** modal, where site-level properties (when any are defined) can be edited. If the site has none, the modal says "This container has no editable properties." View users do not see this button.

The Locations dashboard

Click the **Locations** node in the tree to open the full Locations page. This is the management dashboard for your sites: it adds summary statistics, more table columns, and (for administrators) the **Add** button.

The page header shows the pin icon, the title **Locations**, a subtitle with the location count, and the status dot with the last-update time. Below it, a row of three **summary cards** totals up your site at a glance:

- **Total Locations** — how many locations exist.
- **Normal** — how many are currently normal.
- **Alarms** — how many are in alarm. When the count is above zero the card turns red-bordered and reads "N location(s) in alarm"; otherwise it reads "All locations normal".

Location cards and the code badge

Each location card shows:

- The location's icon and **name**.
- An optional **code badge** — a small monospace blue chip showing the location's short code, for example `SITECODE`. Codes are optional labels you can assign when creating or editing a location; they are handy when your organization refers to sites by an internal code rather than by name.
- The status badge (**Alarm** / **Normal** / **Disabled**) on the right.

- An optional description line.

Disabled locations render faded (at half opacity); locations in alarm get a red border. Clicking anywhere on a card navigates to that location's Groups page.

In **Table** view the columns are **Location | Status | Code | Enabled | Description**. The Code column shows the badge or a dash when no code is set, and the Enabled column shows a green "✓ Yes" or a gray "— No".

Searching the list

Every list page in this chapter has a live search bar above the list. Type into it and the list filters as you type — on this page the search matches the location name, code, and description. An ✕ button appears while text is present; click it to clear the filter. While a filter is active, the page footer shows how many items match, for example "1 of 3 locations".

If the list is empty the page says "No locations found."; if your search excludes everything it says "No locations match the current filter."

The Groups page

Click a location and its Groups page opens, listing the areas or categories defined within that site — for example **Environmentals** and **Power Units** in the Sunbury Office demo site.

The layout mirrors the Locations dashboard: a header (folder icon, title **Groups**, subtitle showing the location name and group count), the Total / Normal / Alarms summary cards ("N group(s) in alarm" / "All groups normal"), the search bar (matching name, status, and description), and the **Cards/Table** toggle.

The **All Groups** section header carries the page's action buttons: **Add** (Admin only), **Properties** (visible to everyone — this is where the *location's* own settings live; see *Editing properties* below), and the view toggle.

- **Card view:** each group card shows a folder icon, the group name, its status badge, and — when a description is set — a description row marked with an ⓘ icon. Disabled groups render faded.
- **Table view:** columns are **Group | Status | Enabled | Description**.

Click a group card or row to open its Units page.

Note: If another user deletes the location you are viewing, the page tells you "Location deleted by another user" and redirects you back to the Locations list after a short countdown. This recovery pattern is described in *Working with the Console*.

The Units pages

Click a group and its Units page opens, listing the monitored devices in that group. The header shows the full path — for example "Sunbury Office > Power Units — 4 Unit(s)" — followed by the familiar summary cards, search bar, and **All Units** section header with **Add** (Admin only), **Properties**, and the view toggle.

Custom data columns

What makes the Units page different is that its columns are configurable per group. A group can define up to five extra columns through its **Header 1 ... Header 5** properties (see **Editing properties** below). Each header names a unit property, and the page shows that property's live value for every unit — so an administrator can surface exactly the readings that matter for that group: temperatures for a sensor group, load and battery figures for a UPS group, and so on.

- **Card view:** each unit card shows the unit's icon and name, its status badge, and an optional ⓘ description row. When custom columns are defined, the card adds a key/value grid below a dashed rule — the column label in uppercase on the left, the live value in monospace on the right. Values the unit has not reported show an italic em-dash.
- **Table view:** columns are **Unit | Status | [your custom columns] | Description**, with the same monospace values and em-dashes for empties.

The search bar on this page matches the unit name, status, description, **and** every custom column value — so you can type a serial number or a reading straight into the search box.

Note: Unlike the Locations and Groups tables, the Units table has no fixed **Enabled** column. If you want one, add "Enabled" as one of the group's headers.

Clicking a unit card or row opens that unit's detail page — live data, alarm settings, and controls — covered in **Units in Detail**.

The environmental card style

Groups that hold temperature/humidity sensors can opt into a special card view that replaces the key/value grid with two large readouts per unit: a red thermometer with the **Temperature** reading and a blue droplet with the **Humidity** reading, each shown in large bold digits.

The readouts are driven by the same custom columns: the page uses the first header whose name contains "temp" and the first containing "humid" (capitalization doesn't matter, so "Temperature", "Temp", and "Rel Humidity" all work). The labels under the digits show your actual header text. If a unit has no matching value — or the group has no matching header — the readout shows an em-dash rather than an error. Any description moves to a centered line beneath the readouts.

Everything else on the page — summary cards, search, the table view (which still shows all custom columns generically), and all the procedures below — behaves exactly as on the standard Units page.

This environmental gauge view is set up by Pythia Technologies during configuration — it is not something you turn on yourself from the Group Properties dialog. Once Pythia has assigned the view to a group, its unit cards read and display the temperature and humidity values automatically.

Managing the hierarchy (Admin)

The rest of this chapter covers building and maintaining the tree. All of these procedures require Administrator access — View users do not see the **Add** buttons, the rename pencil, or the **Save** and **Delete** buttons.

Names for locations and groups follow one rule set everywhere: a name is required, may be 1–100 characters, and must not contain any of the characters \ & % # ? " , . The DM checks the name as you submit the form and shows any problem in red under the field.

Adding a location

1. Open the **Locations** page and click **+ Add**. The **Add Location** dialog opens.

1. Fill in the fields:

Field	Required	Notes
Location Name	Yes	Up to 100 characters; the standard name rules apply.
Enable for monitoring	—	Checked by default. Clear it to create the location disabled.
Code	No	The short site code shown as the blue badge on cards and in the table (up to 250 characters).
Description	No	Up to 200 characters.

1. Click **Save**.

If the name duplicates an existing location, the dialog shows "A location with that name already exists." Other problems appear under the offending field — for example "Location name is required." or "Name cannot contain: \ & % # ? " , ,".

When the save succeeds, the dialog closes, a brief "Saving..." overlay appears while the server finishes setting the location up (see *Working with the Console*), and the tree then navigates to your new location — landing on its empty Groups page, ready for the next step.

Adding a group

1. Open the location's Groups page and click **Add**. The **Add Group** dialog opens.

1. Fill in the fields:

Field	Required	Notes
Group Name	Yes	Up to 100 characters; standard name rules.
Description	No	Up to 200 characters.
Type	—	Generic is the only type available today; additional group types are planned.
Enable for monitoring	—	Checked by default.

1. Click **Save**.

A duplicate name shows "A group with that name already exists under this location." On success the dialog closes, the "Saving..." overlay runs, and the tree navigates to the new group's empty Units page.

Adding a unit

1. Open the group's Units page and click **+ Add**. The **Add New Unit** dialog opens and loads the available networks and drivers — the two dropdowns briefly read "-- Loading networks... --" and "-- Loading drivers... --".

1. Fill in the fields:

- **Unit Name** — the device's display name, up to 100 characters. Unit names cannot contain a backslash, slash, or pipe character.
- **Network** — the network the DM will reach the device on. Each entry shows the network's name and type, for example "LAN (NET)" or "Serial A (COM1)".
- **Driver** — the driver that speaks the device's protocol, also listed as "Name (Type)". When you select a driver, its description appears in italics beneath the dropdown — for example "Generic BACnet/IP Driver" — so you can pick by purpose rather than by name.

1. Click **Save**.

The driver must be compatible with the network you chose: a network-type driver needs a network connection, a serial-type driver needs a serial port. The dropdowns list every driver regardless, and the check happens when you save — an incompatible pairing is rejected with a message such as "Driver type (NET) does not match Network type (COM2)". Pick a matching network or driver and save again.

Other messages you may see: "Unit name is required.", "Please select a network.", "Please select a driver.", or "A unit with that name already exists in this group".

When the save succeeds, the DM creates the unit with sensible protocol defaults already filled in (for example, an SNMP unit starts with community `public` and port 161). After the "Saving..." overlay clears, the tree navigates to the **new unit itself** — not back to the group — so you land directly on the unit's page to finish configuring it. Unit configuration is covered in **Units in Detail**.

Renaming a location or group

Renames live in the Properties modal — there is no rename option on the cards themselves.

1. Open the Properties modal for the item: on a location's Groups page, click **Properties** for the location; on a group's Units page, click **Properties** for the group.
2. Click the pencil icon next to the name at the top of the modal, type the new name, and confirm with the green checkmark (or press Enter). The inline rename flow — checkmark, cancel, Esc, and validation — follows the standard pattern described in **Working with the Console**.
1. A duplicate name is rejected inline: "A location with that name already exists" (or "A group with that name already exists"). Confirming the unchanged name simply cancels the edit.

On success the modal closes and a "Saving..." overlay runs while the server carries the rename through — then the tree navigates to the item under its new name. Everything beneath the renamed item is updated automatically; no links or child pages break.

Editing properties

The same Properties modals hold each item's editable settings.

- **Location Properties** (from the Groups page): typically **Enable**, **Code**, and **Description**.
- **Group Properties** (from the Units page): **Description**, **Enable**, and the **Header 1 ... Header 5** fields that define the group's custom unit columns. Type a unit property name into a header field and, after **Save**, the Units page immediately shows that property as a new column.
- **Site Properties** (from the Site Overview page): site-level fields, when defined.

Edit the fields and click **Save**. Property values may not contain the pipe character — the modal rejects such a value with "Property '<name>' cannot contain the | character." On success the modal closes and the list behind it reloads with the new values.

View users can open every Properties modal but see the fields read-only, with no pencil and no **Save** or **Delete** buttons.

Deleting a location or group

Deletion is strictly bottom-up: a container that still has children cannot be deleted. To remove a location you must first delete its groups, and to remove a group you must first delete its units. (Deleting a unit is done from the unit's own page — see **Units in Detail**.)

1. Open the item's Properties modal and click the red **Delete** button in the lower-left corner.
2. A confirmation dialog asks: "Are you sure you want to delete the location "<name>?" (or "...the group "<name>?") with the caution "This action cannot be undone."
1. Click **Delete** to confirm, or **Cancel** to back out.

Caution: Deletion is permanent. There is no undo and no recycle bin — a deleted location, group, or unit is gone, along with its configuration.

If the item still has children, the dialog changes to a **Cannot Delete** notice instead of deleting anything:

- For a location: "The location "<name>" contains N group(s)." with the instruction "Delete all group(s) before deleting this location."
- For a group: "The group "<name>" contains N unit(s)." with the instruction "Delete all unit(s) before deleting this group."

Click **OK**, remove the children, and try again.

When a delete succeeds, the page you were on no longer exists, so the DM navigates for you: deleting a location returns you to the Locations page; deleting a group returns you to the parent location's Groups page. Anyone else who was viewing the deleted item at that moment sees the countdown redirect described in **Working with the Console**.

What View users see

Everything in this chapter up to **Managing the hierarchy** is fully available with View access: browsing every level, the summary cards, searching and filtering, toggling Cards/Table, manual **Refresh**, and clicking through the tree.

What View users do **not** get:

- The **+ Add** buttons on the Locations, Groups, and Units pages.
- The **Properties** button on the Site Overview page.
- Inside the Properties modals: the rename pencil, **Save**, and **Delete**. The modals themselves open normally, showing all properties read-only — a View user can look up a group's headers or a location's code, just not change them.

Chapter 5 — Units in Detail

Chapter 4 took you as far as the unit — one monitored piece of equipment in the tree. This chapter goes inside it: the unit overview page that shows everything at a glance, the **Live Data** pages where readings live, the **Controls** pages where operators command the equipment, and the unit's alarm settings. Operators will spend most of their day on these pages; administrators use them to configure what gets polled, what gets alarmed, and what can be commanded.

One distinction to fix in mind before anything else: the alarm pages under a unit are **configuration**. They define *what condition to watch* on this unit — the threshold, the severity, how many failed polls it takes to fire. When one of those conditions actually trips, the resulting alarm *event* appears site-wide under the top-level **Alarms** node, alongside events from every other unit. Working those events — acknowledging, filtering, reviewing history across the site — is Chapter 7. This chapter is where you set the rules.

The unit overview page

Click any unit in the tree — say **CRAC-01** under **Sunbury Office > Power Units** — and the unit overview page opens: a single dashboard combining the unit's live readings, its alarm states, and its controls. For most day-to-day monitoring this one page is all you need.

The header and the four-state status block

The header shows the unit's name (with a rack icon) and its description underneath. At the top right sit the familiar status dot and **Updated:** timestamp from Chapter 3, and below them a status block that summarizes the whole unit in one of four states:

Status block	Meaning
ALARM ACTIVE (red, warning icon, with an "N Active" count)	At least one of the unit's alarms is currently active.
COMMUNICATIONS (red, warning icon, no count)	The unit's alarms are all normal, but at least one data point or control failed to update on its last poll.
ALL NORMAL (green, check icon, "0 Active")	The unit is running, every alarm is normal, and every poll is succeeding.
DISABLED (gray, circle-slash icon, no count)	The unit is not running — its Enable setting is off.

The states have a strict precedence: an active alarm always wins over a communications problem, and both are only reported for a unit that is actually running. A disabled unit shows **DISABLED** with no alarm count at all — a stopped unit's counts are stale, so the console deliberately doesn't show them. The unit's name in the page title follows the same logic: blue when normal, red when in alarm, muted gray when disabled.

The toolbar

- **Open Device** — visible to all users, but only when the unit has a URL Open address configured. It opens the device's own built-in web page in a new browser tab, so you can reach the equipment's native interface without hunting for its IP address.
- **Properties** — Admin only. Opens the Unit Properties dialog described at the end of this section.

The Live Data column

The left column, titled **Live Data**, lists every data point on the unit — one row per point with its name, its live value (right-aligned, monospace), and its description. The column title is itself a link: click **Live Data** to jump to the full Live Data list page described in the next section. Clicking any individual point name jumps to that point's detail page.

The coloring tells you about health at a glance: a point's name turns **red** when its last poll failed, and a value turns **red** when the point itself is reporting a problem. On a disabled unit everything renders muted gray instead — a stale reading on a switched-off unit is not an active fault.

If the unit has no points, the column reads "No data points configured for this unit."

The Alarm Status column

The right column, titled **Alarm Status**, shows one card per alarm defined on the unit. Each card carries a colored dot, the alarm's name, and an uppercase status pill:

- **ACTIVE** — red dot and pill, the card tinted dark red with a red left edge. This alarm's condition is currently true.
- **NORMAL** — green dot and pill. Being watched; condition not met.
- **DISABLED** — gray dot and pill. The alarm's Enable setting is off, so it is not being evaluated. (An alarm that is active *and* disabled still shows **ACTIVE** — an active condition always wins.)

Cards sort active-first, then enabled-normal, then disabled, alphabetically within each tier — so anything demanding attention is always at the top. Clicking a card opens that alarm's detail page; clicking the **Alarm Status** title opens the unit's full alarm list.

The Controls section

Units whose driver supports control show a full-width **Controls** table below the two columns (units without controls don't show the section at all). Each row lists the control's name, its **Setpoint** (the commanded value), its **Value** (the live reading from the equipment), and its description. A yellow **OVR** badge next to a name means that control is currently in override. Clicking a control's name opens its detail page; clicking the **Controls** title opens the full controls list.

Refresh and errors

The unit overview refreshes itself every **2 seconds** — it is a live page, and the **Updated:** time should tick along steadily. If the connection to the DM drops, the page keeps showing its last-good content with a

"Reconnecting..." note until the connection returns, exactly as described in Chapter 3. If the unit itself has been deleted by another user, the page shows a plain error panel — this page does not use the countdown redirect.

Unit Properties (Admin)

The **Properties** button opens the standard Properties dialog from Chapter 3: the unit's name at the top with the rename pencil, then the unit's fields — typically Description, Enable, Network, IP, Code, URL Open, and any driver-specific extras. All the Chapter 3 rules apply: rename validates to 1–100 characters with no `\ & % # ? " ,`, a duplicate is refused with "A unit with that name already exists", and a successful save closes the dialog and refreshes the page. Note that some property changes can add or remove the unit's **Data**, **Alarms**, or **Controls** containers in the tree — the tree refreshes automatically after a save.

Delete sits in the dialog footer. Confirming "Delete this unit?" removes the unit *and everything under it*: "This action cannot be undone. All data points, alarms, and controls under this unit will also be removed." After the **"Saving..."** overlay clears, you land on the parent group's page.

Live Data

The Live Data list

Click the unit's **Data** node in the tree (or the **Live Data** column title on the unit overview) to open the full list of data points. The header shows the count ("12 data points") and the page refreshes at the system-wide list rate (30 seconds unless an administrator has changed it).

The table has four columns: **Data | Enable | Value | Description**. Rows appear in the order the driver defines them (a deliberate, meaningful sequence — there is no click-to-sort here). The name coloring follows a simple precedence: a point whose Enable is **No** renders muted gray; an enabled point whose last poll failed renders red; everything else is normal.

When one or more enabled points failed their last poll — and the unit itself is enabled — a red banner appears above the table:

> Communications issue — 3 data points failed to update on the last poll

The banner is suppressed when the unit is disabled, for the same reason the header drops to a plain **DISABLED** pill: stale readings on a stopped unit aren't faults.

Admin only: The toolbar — **Add** and **Properties** — appears only for Administrator accounts. View users see the list with no toolbar at all.

Adding a data point

Click **Add** to open the **Add New Data Point** dialog. The form is driver-specific: alongside the universal fields — the point's name, Description, Enable, Format, and the Gain/Offset/Addend scaling values —

you'll see whatever addressing the driver needs to poll the value, such as a Modbus Function and Register for a Modbus unit, or an OID for an SNMP unit. (Live, read-only fields like the current Value are not part of the form.)

Click **Save** to create the point. Names follow the usual rules; the server refuses a duplicate with "A data point with that name already exists" and rejects names containing backslash, slash, or pipe. On success the dialog closes, the new point appears in the list immediately, and the tree gains the new leaf under **Data**.

The Data Point Properties grid

For configuring points one at a time there is the detail page (next section) — but when you need to sweep a setting across a whole unit, use the bulk grid. Click **Properties** in the toolbar (tooltip: "Configure description, enable, value format and trend for every data point") to open the wide **Data Point Properties** dialog: one row per data point, with editable **Description**, **Enable**, **Value Format**, and **Trend Data** cells.

This grid works differently from the ordinary Properties dialog, and the difference matters:

Note: There is no Save button in this grid. **Every change saves immediately**, the moment you make it — flip a toggle or pick a dropdown value and it is already written. A green "✓ **Saved**" confirmation flashes beside the row for a moment after each successful write. If a write fails, the console tells you which field and which point ('Could not update Enable for "Temperature".') and the grid reverts to the last-known values.

The cells:

- **Description** — free text per point. (Clearing a description doesn't leave it empty for long — the console restores it to the point's name on the next property write.)
- **Enable** — a toggle. Disabled points stop being polled and render muted throughout the console.
- **Value Format** — how the raw value is displayed (the dropdown lists every format the driver offers, plus the point's current format if it isn't in the list).
- **Trend Data** — assigns the point to a trend bucket so the trending service starts recording its history (Chapter 9). The blank choice shows as — **Off** —.

The shaded **Set all** → row at the top applies one value down an entire column in a single stroke — set every point's Enable to On, or every Value Format, or every Trend bucket at once. Pick a value from a Set-all dropdown and it applies to every row, then the dropdown snaps back to "Set all...". Description has no Set-all, since descriptions are individual by nature.

The grid's footer holds a red **Delete** button (when the driver permits it) that removes the *entire* Data container. The confirmation names it explicitly — 'Delete the entire data container "CRAC-01\Data"?' — and warns: "This action cannot be undone. All 12 data point(s) under this unit will be removed." Confirming shows the "**Saving...**" overlay and lands you back on the unit overview.

The data point detail page

Click a data point's name — from the tree, the unit overview, or the Live Data list — to open its detail page: the biggest page in the console, combining the live value, a live chart, exports, and the point's configuration.

The live value

The header shows the point's name, its description as a subtitle, and the current reading as a large blue monospace number. Units of measure come from the description: a description written as `Temperature (F)` displays the big value with **F** beside it and the subtitle as simply "Temperature". (This convention holds throughout the console — put the units in parentheses at the end of the description and the pages and exports pick them up.) The value shows `--` when no reading is available.

If the point's last poll failed, the big value turns red and a red **⚠ COMM ISSUE** badge appears next to the title (its tooltip reads "Last poll failed for this data point").

The page refreshes every **2 seconds** — this is a live page regardless of the system list rate. After the first successful load, a failed poll never wipes the page: the last-good value and chart stay put and only the status dot flags the problem until the next good response.

The live chart

Below the value sits a large chart panel. Understand one thing about it first: **the chart starts empty every time you open the page** ("Collecting data..."), and fills as the page sits open, sampling the live value every 2 seconds. It is a live strip chart of *this viewing session*, not stored history — for recorded history over days and weeks, assign the point to a Trend bucket and use the Trends pages (Chapter 9).

Choose the presentation with the **Chart:** dropdown:

- **Line** (the default) — a blue line with an area fill, auto-scaled value axis, clock-time labels across the bottom, and a dot marking the newest sample at the right edge. Samples drift left as time passes and fall off the end of the window.
- **Bar** — the same axes with one vertical bar per sample; suited to count-style readings.
- **Gauge** — a 180° dial showing the instantaneous value against colored zones (see below). No time axis.
- **Step** — offered only for two-state (digital) points such as On/Off or Open/Closed formats. The trace holds each state as a horizontal level and jumps vertically at the moment of change, with the two axis labels taken from the format's own words (Off/On, No/Yes, Open/Closed). The last state extends to the right edge — "still in this state now".

In Line, Bar, and Step modes a **Window:** dropdown selects how much time the chart spans: **Last 1 minute**, **Last 5 minutes** (the default), or **Last 15 minutes**.

Hover the mouse anywhere over a Line, Bar, or Step chart and a dashed cursor with a highlight dot appears, plus a small box showing the exact sample: the value with its units (or the state word for a

digital point) and its HH:MM:SS timestamp.

Note: Your chart choice is remembered **per data point** on this browser — pick Gauge for the temperature point and Line for the humidity point, and each page reopens the way you left it. The time window is remembered the same way.

The gauge and its display zones

In Gauge mode the Window control is replaced by four small numeric inputs: **Min**, **Max**, **Warn Low**, and **Warn High**. They shape the dial's colored arc: red caps at both extreme ends, yellow between Min→Warn Low and Warn High→Max, and green between the two warn points, with tick labels at all four values. The console pre-fills sensible values derived from the current reading; type your own into any field and that field sticks (it stops auto-adjusting and is remembered for this point on this browser). Clear a field to let the console derive it again.

Caution: The gauge's Min / Max / Warn Low / Warn High values are **display-only**. They color the dial on your screen and do nothing else — they are stored only in your browser, they are not sent to the DM, and they do **not** trigger alarms or notifications. The real thresholds that fire alarms are configured in the unit's alarm settings, later in this chapter. Do not set a Warn High on a gauge and assume anyone will be notified when the needle crosses it.

The Trend selector

If the point supports trending, administrators see a **Trend:** dropdown in the controls row — **Off** plus the available trend buckets. Changing it saves immediately, no dialog needed; choosing a bucket starts the trending service recording this point's history. This is the same setting as the **Trend Data** column in the bulk grid — use whichever is closer to hand.

Exporting the chart and its data

The **Export:** buttons are available to all users:

- **PDF** — a landscape page with the point's name, description, the export date and time, and a snapshot image of the chart exactly as currently displayed. Saved as `<name>.pdf`.
- **Excel** — a spreadsheet file (`<name>.xls`) with two columns, `Timestamp` and `Value (units)`, one row per sample currently in the chart's buffer.

Since the chart buffers only what it has collected while open, an export covers at most the selected window of the current viewing session — let the page sit open for the span you want before exporting.

Data point Properties (Admin)

The **Properties** button opens the point's full Properties dialog: the name with its rename pencil at the top, then Description, Enable, Format, the Gain/Offset/Addend scaling values, Trend, and the driver's addressing fields. All the Chapter 3 conventions apply. Two behaviors specific to data points are worth

knowing: clearing Gain, Offset, Addend, or Description doesn't delete them — the console restores the defaults (Gain 1, Offset 0, Addend 0, Description = the point's name); and after a save the page refreshes in place *without* discarding the chart you've been collecting.

Delete (in the dialog footer) asks 'Delete Temperature?' with "This action cannot be undone." — confirming removes the point and returns you to the Live Data list.

Controls

Controls are where the DM stops being read-only. A control is a writable point on the equipment — an on/off command, a numeric setpoint — and the pages in this section both display it and command it.

Caution: Control commands are sent to the equipment without a confirmation step. Clicking **On** or **Off**, or clicking **Apply** on a setpoint (or pressing Enter in the setpoint box), sends the command immediately — there is no "Are you sure?" dialog between your click and the equipment. The DM then pushes the value to the device on its next poll and keeps re-asserting it whenever the device's actual value differs from the commanded one. Treat every click in a Command panel, and every setpoint cell in the bulk grid, as a live action on real equipment.

The Controls list

Click the unit's **Controls** node in the tree (or the **Controls** title on the unit overview) to open the list. The header shows the count ("4 controls") and the page refreshes at the system list rate.

The table has six columns:

Column	What it shows
Control	The control's name — a link to its detail page. Renders red when its last poll failed.
Enable	Yes or No .
Status	Normal (green), Fault (red — enabled but the last poll failed), or Disabled (muted — Enable is off). Disabled takes precedence over Fault.
Setpoint	The commanded value (monospace; an em-dash when none is set).
Value	The live reading back from the equipment.
Description	Free text.

As on the Live Data list, a red banner appears above the table when polls are failing: "Communications issue — 2 control(s) failed to update on the last poll".

Admin only: The **Add** and **Properties** toolbar buttons, and everything in the rest of this section, are Administrator features. View users see the list and the read-back values only.

Adding a control

Add opens the **Add New Control** dialog, which works in two steps. Step 1 asks for the control's name and — for drivers that offer more than one kind — its **Type** (for example a plain setpoint, an on/off control, or an on/off control with override). Once you pick a Type, step 2 appears below a divider with that type's own fields, which vary by driver. **Save** stays disabled until a Type is chosen. On success the dialog closes and the new control appears in the list immediately.

The Control Properties grid

The toolbar's **Properties** button (tooltip: "Configure enable, setpoint and description for every control") opens the **Control Properties** grid — the controls counterpart of the data grid, with one row per control: **Enable** toggle, **Setpoint**, read-only **Value**, and **Description**. The setpoint cell adapts to the control's kind: on/off controls get a three-choice dropdown (blank / On / Off); numeric controls get a text input.

Like the data grid, **every edit saves immediately** with a "✓ Saved" flash — and remember, for the Setpoint column "saves immediately" means *commands the equipment*. There is no confirmation on individual cells.

The Set-all row at the top offers two things:

- An Enable **Set all...** picker (On / Off) that applies down the column.
- A **Clear Setpoints** button — the one place in the controls area that *does* confirm first: "Clear the setpoint on all 4 control(s)? Each is re-read fresh from the unit on the next poll." Confirming blanks every stored setpoint; on the next poll the DM adopts each control's *actual* reading as its new setpoint. This is the "resync to reality" tool — use it when stored setpoints have drifted from how the equipment is really configured and you want the DM to stop pushing old values.

The footer's red **Delete** button (when the driver permits it) removes the whole Controls container, with the usual named confirmation and "cannot be undone" caution.

The control detail page

Click a control's name to open its detail page. The header shows the name and, in large monospace, the control's live value; a red **Comm Issue** badge and red value flag a failed poll. This page polls every **2 seconds**, so you see command results almost as they happen.

The Command panel

Controls with a setpoint show a **COMMAND** section — the operator surface:

- The readout, visible to all users: **Commanded** (blue — the setpoint stored in the DM) with an arrow to **Actual** (white — the live value read back from the equipment). When the two match, the equipment is doing what it was told.
- The command controls, Admin only:

- **On/Off controls** show segmented buttons — for example **On | Off** — with the active state filled blue. **One click sends the command.**
- **Numeric controls** show a monospace input with a blue **Apply** button; Enter in the input also applies. (The 2-second refresh never overwrites a value you are in the middle of typing.)
- Controls of the on/off-with-override type show a separate **Override** row with the same widget style, for forcing the output independently of the normal setpoint.

A feedback line under the controls reports each command and clears itself after a couple of seconds: green on success ("Setpoint → ON sent."), red with the server's message — or "Command failed." — on error.

After a successful send, watch the readout: **Commanded** shows your new value at once, and **Actual** catches up within a poll cycle or two as the DM pushes the command to the device. If Actual never converges, check the control's status and the unit's communications state.

Properties and delete

Below the Command panel, a read-only **Property | Value | Description** table shows the control's remaining configuration (the setpoint and override live in the Command panel, so they're not repeated here). Admins get a **Properties** button opening the standard dialog — rename pencil, editable fields, and a red **Delete** in the footer ("Delete control "Fan Speed"? / "This action cannot be undone."). Deleting returns you to the Controls list.

Unit alarm settings

This is where you tell the DM what to watch on a unit. Each entry here is an alarm **definition**: a condition (typically a comparison against a threshold, addressed the same way the driver addresses data points), a classification (Category and Severity), a persistence requirement (Fail Count), and an Enable flag. To repeat the rule from the top of the chapter: when a definition trips, the **event** it generates is worked from the site-wide **Alarms** pages in Chapter 7 — the pages here define and tune, they are not where you acknowledge or review site activity.

The alarm list

Click the unit's **Alarms** node in the tree (or the **Alarm Status** column title on the unit overview) to open the list, headed **Alarm Status** with a count of definitions.

Five columns: **Alarm | Enable | State | Category | Description**. The **State** pill shows each definition's live evaluation:

- **ACTIVE** (red) — the condition is currently true.
- **NORMAL** (green) — being evaluated; condition not met.
- **UNKNOWN** (gray) — not being evaluated, typically because the definition is disabled.

Active definitions sort to the top, and this page's status dot doubles as a summary: it turns red whenever **any** listed alarm is active, green otherwise. Clicking an alarm's name opens its detail page.

Admin only: The **Add** and **Properties** toolbar buttons and everything below are Administrator features.

Adding an alarm

Add opens the **Add New Alarm** dialog — a single flat form (no type picker) of every field the driver's alarm schema defines. Typical fields: the alarm's name, **Severity**, **Count**, **Category**, **Description**, and then the driver-specific **trigger** fields that define the condition — for a Modbus unit, for example, a Modbus Function and Register to read, an **Operator**, and a **Threshold** to compare against. On success the dialog closes, the list refreshes, and the tree selects the new definition under **Alarms**.

The Configure Alarms grid

The toolbar's **Properties** button (tooltip: "Configure category, severity, fail count, enable and the global alarm test for every alarm") opens the wide **Alarm Properties** grid — the bulk classification editor for every definition on the unit at once.

At the top, above the grid, sits the **Global Alarm Test** toggle. Switching it on pushes the test flag down to **every** alarm on the unit at once — the way to exercise the whole unit's alarm and notification chain without waiting for real conditions. Like everything in these grids it saves immediately, with a "✓ Saved" flash.

The grid itself has one row per definition:

- **Category** — a dropdown for grouping and policy: — **None** — plus the driver's category list (or the standard **Category 01** through **Category 25** when the driver doesn't define its own).
- **Severity** — — **None** — plus the driver's list, or the standard **INFORMATION / WARNING / MINOR / MAJOR / CRITICAL**. Severity drives the color coding on the alarm detail page and the handling of events on the site-wide Alarms pages.
- **Fail Count** — a **1–5** picker: how many **consecutive** polls the condition must hold before the alarm fires. Raise it to ride out momentary blips on a noisy reading; leave it at 1 for conditions that must fire instantly.
- **Enable** — a toggle. A disabled definition stops being evaluated entirely and shows the gray **UNKNOWN** pill in the list.

Every cell saves the moment you change it, exactly as in the other grids. The **Set all...** row applies a value down a whole column — but unlike the data and control grids, each Set-all here **asks for confirmation first**, since reclassifying every alarm on a unit is a bigger decision: "Are you sure you want to set severity to 'CRITICAL' for all 8 alarms?", "Are you sure you want to disable all 8 alarms?", and so on. If some rows fail to update, the console reports the count ("2 of 8 alarm(s) failed to update.").

The footer's red **Delete** removes the entire Alarms container, with the usual named confirmation and count caution.

The alarm detail page

Click a definition's name to open its detail page — the live state of this one condition, plus its complete occurrence history.

The big status word under the title tells you the state at a glance. At rest it reads **NORMAL** in green (or **UNKNOWN** in gray when the definition isn't being evaluated). While the alarm is firing, the display shows the alarm's **severity word** — **CRITICAL**, **MAJOR**, **MINOR**, **WARNING**, or **INFORMATION** — and both the word and the page title take the severity's color: red for CRITICAL, orange for MAJOR, yellow for MINOR, teal for WARNING, light blue for INFORMATION. The state display polls every 2 seconds, so it flips within moments of the condition changing.

The history table

Below the state sits this alarm's own occurrence history: **Time Start | Time Stop | Duration | Text**, newest first (click a column header to re-sort; click again to flip direction). While the alarm is firing, a live row sits pinned at the top — red-tinted, with a bold start time, a **blank** Time Stop, and a **Duration that ticks up once per second** like a stopwatch. When the condition clears, the row closes with its stop time and joins the history. Durations format compactly: 45s , 2m 3s , 1h 2m 3s , 1d 2h 3m .

The history here is scoped to this one definition on this one unit. The same occurrences also appear in the site-wide history in Chapter 7, mixed with every other alarm on the site — this page is the place to study one condition's behavior over time.

Exports

The **PDF** and **Excel** buttons above the table are available to all users. The PDF is a branded landscape report — "Alarm History — <name>", timestamp, record count, and the striped table; the Excel export is a workbook of the same rows. Both include a currently-active occurrence, with its Time Stop exported as "(active)".

Alarm Properties (Admin)

The **Properties** button opens the definition's full dialog: the name with its rename pencil, then every field — including the **trigger fields** (Operator, Threshold, and the driver's addressing) that the bulk grid does **not** expose. This dialog is where the actual alarm thresholds live — the real ones, as opposed to the display-only gauge zones from earlier in this chapter. Severity, Category, and Count can be edited here too, individually; the bulk grid does the same across all definitions at once.

Delete in the footer removes the definition ('Delete High Temperature?' / "This action cannot be undone.") and returns you to the alarm list. Deleting a definition does **not** delete its recorded history — the closed occurrences stay in the database; only **Clear History** removes them.

Clear History (Admin)

The red **Clear History** button above the table permanently deletes this alarm's closed occurrences. It is disabled when there is nothing to clear, and confirms first: "Permanently delete 47 record(s) for this alarm?" with "This action cannot be undone."

Caution: Clear History cannot be undone — the records are deleted from the database, and they disappear from the site-wide history too. Export to PDF or Excel first if you need to keep the record. An occurrence that is **currently active** is not affected; it remains as the live row and will write its own history entry when it closes.

What View users see

Everything in this chapter that displays information works for View users: the unit overview with all four header states (including **Open Device**, when configured), the Live Data list, the data point detail page with its charts, gauge inputs, and PDF/Excel exports, the Controls list and the **Commanded** → **Actual** readout on the control detail page, the alarm list, and the alarm detail page with its history and exports. What View users never see are the write controls: no **Add**, **Properties**, or **Clear History** buttons, no bulk grids, no **Trend** selector, and — most importantly — no command buttons in any Command panel. The buttons are absent, not grayed out, and the DM refuses any write attempt from a View account regardless.

Chapter 6 — Equipment Dashboards

When a unit is set up with a matching device driver, the DM replaces the generic unit page from Chapter 5 with a custom **equipment dashboard**: a live, animated diagram of that specific piece of equipment, drawn from its real data. A UPS shows its power path, a generator shows its engine and transfer switch, a battery monitor shows every cell. This chapter first explains the behaviors every dashboard shares, then walks through each dashboard in the gallery so you know what to expect from your own equipment.

What an equipment dashboard is

Not every unit gets one. Units without a matching driver use the standard unit page described in Chapter 5 — same information, generic layout. Units *with* a matching driver (a supported UPS, generator, CRAC, battery monitor, power distribution unit, or temperature/humidity sensor) open a dashboard tailored to that model instead. You do not choose between them; the console opens the right page automatically when you click the unit in the navigation tree.

Every dashboard keeps the familiar unit-page frame:

- **Header** — the unit's name and description, the four-state status pill (**ALARM ACTIVE** / **COMMUNICATIONS** / **ALL NORMAL** with an active-alarm count / **DISABLED**), the status dot, and the **Updated: HH:MM:SS** clock.
- **Toolbar** — **Open Device**, which opens the equipment's own built-in web page in a new tab (shown only when the unit has a device address configured; available to all users), and **Properties** for editing the unit's settings.
- **The equipment diagram** — the centerpiece, described below.
- **Live Data** — every data point the unit reports, with its current value and unit label. On most dashboards the points are grouped to mirror the diagram (Input, Battery, Output, and so on). Every row is clickable: click a reading to jump to that data point's detail page with its live chart.
- **Alarm Status** — one row per alarm, sorted so you never hunt: active alarms first (red left border and tinted background), enabled-but-normal alarms in the middle, disabled alarms grayed at the bottom. Click any row to open that alarm's detail page.

The **Live Data** and **Alarm Status** section titles are themselves links — click a title to open the unit's full data or alarms list in the tree.

Admin only: The **Properties** button appears only for Administrator accounts. It opens the same Properties dialog described in Chapter 3, including the rename pencil and the red **Delete Unit** button. Deleting from here is a cascade delete — "This will permanently delete the unit and every Data point, Alarm, and Control under it" — and cannot be undone. View users see the identical live dashboard without the **Properties** button.

Note: Equipment dashboards are for watching, not commanding. No dashboard sends control commands to the equipment. Where setpoints exist (the CRAC dashboard), they are shown read-only; changes are made on the unit's **Controls** page (Chapter 5).

How every dashboard behaves

The dashboards look different but run on the same machinery. Learn it once here.

Live refresh

Every equipment dashboard refreshes automatically every **2 seconds** — the footer reads "Auto-refresh: 2 seconds", and the **Updated** clock in the header advances with each cycle. While an administrator has the **Properties** dialog open, refresh pauses so live updates do not fight the edit; it resumes the moment the dialog closes.

Mode painting and flow animation

On each refresh, the dashboard reads the equipment's reported operating state and repaints the diagram to match:

- **The active path lights up and flows.** Power lines, air ducts, and bus bars that are carrying load show a moving-dash animation in the direction of flow. On the UPS dashboards, the battery line even reverses direction — dashes flow **into** the battery while charging and **out of** it while discharging.
- **Inactive paths dim.** A bypass rail that is not in use, a generator that is on standby, a phase leg with no voltage — these render muted so the live path stands out.
- **A mode badge shows the state in words.** Each dashboard has a mode display — on the UPS pages it is styled as the unit's own front-panel LCD — reading **ONLINE**, **ON BATTERY**, **RUNNING**, **FLOAT**, and so on, with a chassis status LED that tracks the mode color.

Fault highlighting

After the mode is painted, active alarms are overlaid onto the picture:

- A **fault** (serious alarm) paints its component red with a pulsing glow, and any flow through it freezes — a frozen red path reads as broken.
- A **warning** paints its component steady orange while flow continues — degraded, but still working.
- Multiple active alarms combine: a battery fault and an inverter warning each light their own part of the diagram at the same time.
- Alarms that do not map to one visible component (an emergency-power-off contact, a fan failure, an overtemperature) light a pulsing **warning triangle** next to the mode badge. Hover the triangle to see the names of the active faults in a tooltip.

When communications are lost

If the DM's last poll of the device failed, the dashboard makes it unmistakable that you are no longer looking at live data:

- The header pill turns **red**.
- Every affected reading — in the diagram's data panels and in the **Live Data** list — is replaced with dashes (---). Dashes mean "no current value," never zero.
- A **Communications** alarm goes active in the **Alarm Status** list, so the outage is recorded and can drive notifications like any other alarm (Chapter 8).

Normal live display resumes automatically on the first successful poll after the device answers again.

Note: A **DISABLED** unit is different from a communications loss. A disabled unit is not being polled at all, so its dashboard shows the gray **DISABLED** state rather than a communications failure.

Dashboard gallery

The rest of this chapter is a tour of each dashboard: what the diagram shows and what is unique to that equipment. Skip to the sections that match your site.

Note: Operating-state and alarm names on these dashboards come from the equipment itself and are shown exactly as the device reports them — including the manufacturer's own spellings, such as "Genrator in Off" on the Generac controller.¹

Vertiv GXT5 UPS — the reference dashboard

The GXT5 dashboard (unit **UPS-21A** in the figures) is the pattern all the others follow. The diagram is a cutaway of the rack-mount chassis showing the full double-conversion power path: **AC IN** → **RECTIFIER** → **DC BUS** → **INVERTER** → **STATIC SWITCH** → **LOAD**, with the battery hanging off the DC bus and a dashed **STATIC BYPASS** rail across the top. Analog meter icons mark the input and output feeds.

What to watch:

- **The UPS Mode display**, drawn as the unit's front-panel LCD, reads **ONLINE**, **ON BATTERY**, **BYPASS**, or **DISABLED**. Online, the main chain flows green and the battery line trickles downward (charging). On battery, the battery-to-inverter path turns orange and the battery line reverses upward (discharging). On bypass, the amber bypass rail flows while the main chain dims. If the rectifier or charger fails while the load stays up, the battery-fed path turns yellow — still working, but on a degraded source.
- **Three live panels** — **Input** (voltage, current, frequency, power factor), **Battery** (voltage, charge %, runtime, temperature), and **Output** (voltage, current, load %, VA, watts) — update every 2 seconds.
- **The Battery Status badge** above the Live Data list summarizes the battery at a glance: **Fully Charged**, blue **Charging**, orange **Discharging**, or a pulsing red **Low Battery**.

- **Live Data is grouped** into Input / Bypass / Battery / Output / Unit sections whose title colors match the diagram panels, so a reading is easy to trace back to the picture.

Roughly thirty of the UPS's alarms are wired to specific diagram components — a battery fault reddens the battery, an inverter fault reddens the inverter — with the rest lighting the warning triangle.

Vertiv EXM 3-phase UPS

The EXM dashboard is the GXT5 scaled up to a three-phase, modular, floor-standing UPS.

Beyond the GXT5 pattern, it adds:

- **Two bypass rails** — a **STATIC BYPASS** and an outer **MAINTENANCE BYPASS** that wraps the whole unit — with an extra **MAINT BYPASS** mode when the unit reports "Load On Maintenance Bypass" (the EXM's own spelling¹).
- **A power-module bay** of five drawn slots; the live "Active Power Modules" count lights that many slots green, and module alarms (fan, fuse, temperature, power supply) highlight the bay.
- **An external battery cabinet** below the DC bus, matching the EXM's real construction, with per-cabinet disconnect alarms.
- **Per-phase data grids** — Input A/B/C, Bypass A/B/C, Output X/Y/Z with totals — instead of single-phase panels.
- **A GENERATOR overlay** — when the UPS reports it is being fed by a generator, the AC-input feed turns amber and is labeled **GENERATOR**.

Its alarm map also covers parallel-system faults (parallel cable, load sharing, loss of redundancy) for multi-unit installations.

Battery string monitor

The battery dashboard (unit **Battery System**) monitors two battery strings cell by cell.

- **Each string is a live grid of 24 cells** (two rows of twelve), colored from that cell's own voltage and temperature readings. Both strings feed a DC bus bar that flows toward **TO LOAD** while the batteries are discharging.
- **Weak cells stand out automatically.** Every cell is compared with its string's average: a cell drifting moderately from the average outlines **amber**; one drifting badly outlines **red**. Hover any cell for a tooltip with its live voltage and temperature; click a cell to jump straight to its voltage data point.
- **A per-string alarm-count badge** shows how many alarms each string currently has.
- **The mode display** reads **FLOAT**, **DISCHARGING**, **EQUALIZING**, **STRING OFFLINE**, or **THERMAL RUNAWAY**.
- Where the monitor raises an alarm without saying which string it belongs to, the dashboard highlights **both** string panels — the cell grid itself then shows you which cell is the culprit.

Live Data is grouped **String 1 / String 2 / System**.

Generac generator

The generator dashboard (unit **Genset 1**) is the most animated page in the console: a full genset scene with a radiator whose **fan spins** (fast when running, slow while cranking), an engine block that **puffs exhaust smoke** while running, the alternator, the starter battery, the fuel-tank base, and the H100 control panel — plus an **automatic transfer switch (ATS) whose blade physically swings** between **UTILITY** and **GENERATOR** to show which source is feeding the load.

- **Modes** cover the whole start sequence and its failures: **STANDBY** (engine gray, utility feeding the load), **CRANKING** (amber, fan turning slowly, no output yet), **RUNNING** — with badge variants **EXERCISING**, **WARMING UP**, **COOL DOWN**, and **RUNNING • ON GEN** when the ATS has transferred — **SHUTDOWN** (emergency stop or overcrank; the engine pulses red while utility keeps the load), and **POWER LOSS** (generator on standby *and* utility lost — the site is dark, painted red).
- **The NOT-IN-AUTO warning strip** appears below the diagram whenever the generator cannot start on its own — the controller reports flags such as "Generator Not in Auto", "Switch In Manual", or "Genrator in Off"¹. The strip lists the exact flags and tints the control panel amber. For a standby generator this is one of the most important things the dashboard can tell you: **a generator that is not in Auto will not start when the power fails.**

Live engine readings — RPM, coolant temperature and level, oil pressure and temperature, output voltage/current/kW/frequency, starter-battery voltage — are placed on the diagram next to the parts they describe.

Liebert CRAC (precision cooling)

The CRAC dashboard (unit **CRAC-01**) draws the air handler in **cross-section** so you can follow the air path: warm **RETURN** air (amber, animated) enters, passes the **FILTER**, crosses the **DX COOLING COIL** (which lights blue and animates while actively cooling), the **HUMIDIFIER** (steam wisps appear while humidifying), and the spinning **BLOWER**, then leaves as cold blue **SUPPLY** air. The refrigeration circuit — **COMP 1**, **COMP 2**, and the **GLYCOOL ECON** economizer — is drawn below the cabinet.

- **A readout strip** under the diagram shows **Temperature**, **Humidity**, **Capacity**, and **Stages** tiles; a tile turns red when its related alarm is active.
- **Alarms pin badges onto the picture** — a bolt, water drop, filter, or fan glyph appears next to the affected zone — and the most serious conditions (smoke detected, loss of power) throw a full-cabinet red overlay and stop the airflow animation. Every active alarm also appears as a chip in a strip below the diagram, so alarms with no natural place on the picture are never hidden.
- **Readings are grouped into cards**: Environment, Operating Status (On/Off pills for Unit On, Cooling, Humidifying, and the like), Run Hours (compressors, main fan, humidifier), and Additional Readings.
- **This is the only dashboard with a Setpoints table** — each control's live value beside its configured setpoint. The table is read-only here; change setpoints on the unit's **Controls** page (Chapter 5).

Power distribution one-line

The power distribution dashboard shows a 3-phase power center as a classic **one-line diagram**: input legs **A/B/C** into the transformer (drawn with its coils and core), output legs **X/Y/Z** to the load, the dashed neutral, a ground symbol, and a bypass arc that appears only when the unit reports it is on bypass.

- **Each phase leg is painted independently.** A leg animates only while its voltage reading is present and above zero; a dead leg renders gray, dashed, and frozen with its voltage dimmed — even while the rest of the diagram is in alarm. One glance tells you which phases are live.
- **The readout strip** shows kVA, kW, frequency, and **Load %** (the highest-loaded output phase).
- **Alarms land on the component they describe:** transformer overtemperature badges the transformer, output over/under-voltage puts arrows on the output legs, an overload puts a bolt on the load leg, ground problems light the ground symbol, and an emergency shutdown covers the whole diagram with a red overlay. Contact-type alarms with no diagram home appear as chips below the picture.

Readings are organized per phase; anything the diagram does not place appears under **Additional Readings**.

Temperature/humidity gauges

Simple sensors get a dashboard too. Any temperature/humidity unit (such as **Office TH1** or **Panel TH1**) shows two large semicircular gauges — **Temperature** and **Humidity** — each with a white needle, a big live value, a status pill, and its alarm thresholds labeled underneath.

- **The colored zones are your real alarm thresholds.** The orange low zone, green normal band, and red high zone are computed from the unit's actual configured alarms (Chapter 5) — so if the High Temperature alarm is set at 70 °F, the red zone starts at exactly 70 °F, and the **Low Alarm / High Alarm** labels under the gauge show the same numbers. Change the alarm thresholds and the gauge zones follow. (This is different from the display-only gauge zones on a data point's detail page — see Chapter 5.)
- **The pill under each gauge** reads **NORMAL**, **HIGH TEMP**, **LOW TEMP**, **HIGH HUMIDITY**, or **LOW HUMIDITY**; while in alarm, the big value pulses red. Click the value to open that data point's detail page.
- Below the gauges sit the same **Live Data** and **Alarm Status** columns as every other dashboard. This page refreshes every **2.5 seconds**.

Because it is generic, the same gauge dashboard serves any make of sensor at your site.

Fleet views

Two dashboards look across **many** devices at once rather than into one. They live under **Views** in the navigation tree, refresh every **5 seconds**, and share the Cards/Table view switch from Chapter 3 (your

choice is remembered).

Branch Circuit Monitor rollup

The **Branch Circuits** view rolls up an entire branch-circuit-monitoring tree — customers, racks, circuits, panels, down to individual breakers — using one page at every level. The title, summary numbers, and columns adapt to whatever level you are viewing.

- **Summary squares** across the top: total items at this level (with active/disabled counts), active items, alarms, and total kW.
- **Each card** shows the item's name, status pill, description, its big live **kW** (or current, at the breaker level), and a **min/max range bar** with a live position marker so you can see where the present load sits in its historical span. Table view shows the same fleet as rows.
- **Click a card to drill one level deeper** in the same page; breaker leaves open their properties page.
- **The chart icon on each card opens a live trend** — a rolling kW-or-current line sampled every 2 seconds over a 5-minute window. The chart starts collecting when you open it; it is not historical data.

Admin only: Administrators can add a new item at the current level, edit each child's enable state and description via the gear icon, and delete items. Delete here is strict — an item that still contains children is refused, and the confirmation dialog says so. (Contrast with the cascade delete on unit dashboards.)

Environmental Sensors wall

The **Sensors** view lists every environmental sensor at the site on one wall, sorted by description — the page to put on a big screen for the whole room's temperatures.

- **Summary squares:** total sensors (active/disabled), active alarms, and the mid/min/max temperature across all enabled sensors.
- **Each card is a thermometer gauge** — a red mercury column on a fixed 32–120 °F scale — with the big live temperature and a status pill. Table view lists sensor, enable state, status, temperature, and description.
- **Click a sensor to open a live rolling temperature chart** — sampled every 2 seconds over a 5-minute window. These sensors are not in the trend database, so there is no history from before you opened the chart; the dialog notes this.
- **PDF and Excel export buttons** export the sensor table — unique to this view among the dashboards.

Admin only: Administrators can edit each sensor's enable state and description via the gear icon. Sensors cannot be added or deleted from this page — they come from the sensor collector automatically. View users see the same live wall and charts without the gear icons.

¹ Status text is shown exactly as the device reports it.

Chapter 6 (continued) — New Equipment Dashboards

These subsections extend the **Dashboard gallery** earlier in this chapter. They cover additional equipment dashboards: cooling units, pumps, the manager and group rollups that watch a whole plant, and the floor-plan views. Every dashboard here shares the frame, live refresh, mode painting, fault highlighting, and communications-loss behavior described earlier — these entries describe only what is unique to each, so read "How every dashboard behaves" first if you skipped ahead. Throughout, **SITECODE** stands in for your own site.

Cooling equipment

Liebert Mini-Mate2 ceiling unit

The Mini-Mate2 dashboard covers a ducted split cooling unit suspended above the drop ceiling, drawn in cross-section: hanger rods drop from the slab, the unit sits above the tile line, a return grille is on the left and a supply diffuser on the right, and the conditioned space is below. It follows the Liebert CRAC pattern scaled down to a single compressor with a GlyCool economizer in the plenum row; the cooling coil and compressor light only when the unit reports cooling capacity above zero. Conditions with no place on the picture — a standby-unit indication or a local alarm — appear as chips beneath the diagram.

Trane Voyager rooftop unit (RTU)

This dashboard draws a packaged rooftop unit on its roof curb: an outside-air hood, then filter, DX cooling coil, electric-heat section, and blower across the cabinet, with the two-compressor condenser end and return and supply air passing through the roof deck. Supply air turns blue while a cooling stage runs and orange while heating; the blower and airflow animate only when the unit is on and its supply fan is running. Alarms the unit flags as "Hand" (a component switched to manual at the equipment) paint amber, while "Fail" alarms paint red with a bolt or fan badge.

Admin only: Unlike every other dashboard, the RTU's Controls section is editable in place for Administrators — **Return Air Setpoint** and **Supply Air Static Setpoint** as numeric fields, and a **Run Mode Selector** drop-down (Automatic, Manual Occ Run, Manual Unocc Run, Manual Stop). Leaving a field empty releases that setpoint back to the equipment's own program. View users see the same rows read-only.

Liebert DS precision cooling unit

This dashboard draws a floor-standing downflow DX cooling unit in cross-section: warm return air enters the top, passes the filter, DX cooling coil, and humidifier/reheat element, then dual blowers push

cold supply air down into the raised-floor plenum and out both directions. Live return and supply temperatures sit on the diagram, and the coil, compressors, free-cooling box, reheat, and humidifier light as the unit reports each mode. Because this Liebert reports a very large alarm set, the **Alarm Status** list has its own scroll bar so it never dwarfs the data panels; the diagram-mappable faults still badge their component, and the most serious conditions (smoke, loss of power, fire) throw a full-cabinet overlay. The same page serves several identical units at a site (for example AHU-02 and AHU-03).

Liebert dry cooler

The dry-cooler dashboard draws a long outdoor cabinet with eight fans in two interleaved stages (Stage 1 is the outboard fan pairs, Stage 2 the center four). Running fans spin and throw heat plumes above their bank, a glycol stub shows animated loop flow, and a stage fault rings that bank red and drops a bolt badge. Every dry cooler at a site shares this one page.

Pumps

Glycol-loop pump

The pump dashboard shows a skid-mounted glycol pump driven by a combination motor starter: insulated suction and discharge piping, the motor and volute with a spinning impeller on a red skid, and the starter cabinet with its RUN lamp, HAND-OFF-AUTO dial, and command LED. Because the pump reports only a commanded state, a run-feedback bit, and one fail alarm, the dashboard's whole job is making disagreement loud — the state pill over the diagram reads **PUMP FAIL** (red), **COMMANDED ON — NOT RUNNING** (amber), or **RUNNING WITHOUT COMMAND** (amber, which usually means someone left the starter in HAND). One page serves every pump at the site.

Group rollups

These pages watch a whole set of like units on one screen and live under a **group** container in the navigation tree. They refresh every **5 seconds** and classify their children by data shape, so adding or renaming a unit needs no page change.

Dry cooler group

The Dry Coolers group page rolls the glycol-cooling plant onto one screen: a row of live cooler cards — each a mini two-fan cabinet with **Stage 1** and **Stage 2** pills that spin and redden with the unit's real state — over a system-status panel. The panel is the "Dry Cooler Manager" (loop supply and return temperatures plus its summary alarms), found automatically by its data shape rather than by name. Group-wide active alarms list as chips ("`<unit>: <alarm>`"), and clicking any card opens that cooler's own dashboard.

Pumps group

The Pumps group page shows the glycol loop as a live diagram: a closed supply/return loop with one branch per real pump — motor, volute with spinning impeller, isolation valves, and animated flow while

running — and a status pill under each (**RUN**, **OFF**, **MISMATCH**, **ALARM**, or **DISABLED**). A **System Status** panel below rolls up the "Pump Manager" points and alarms, its title linking through to the manager dashboard. Click any pump branch to open that pump's page.

AC room group

The AC room page watches a room cooled by a lead/standby air-handler team, and the same page is stamped on several rooms unchanged. Each air handler is a card with a fan glyph that spins when the unit is on, live temperature and humidity, an **ON/OFF/ALARM** pill, and a gold **LEAD** or gray **STANDBY** badge drawn from the team manager's current lead selection; a disabled unit gets a gray **DISABLED** pill and its role badge is hidden. A **Team Manager** panel shows lead, last changeover, and failover tiles and links through to the manager dashboard. On a room that has no manager, the page simply renders every unit as a card.

Aggregate / manager views

These dashboards watch a *manager* — a coordinating unit that runs a set of equipment — rather than a single machine.

AC lead/standby team manager

This is a control diagram, not a picture of a machine: it shows a Pythia manager that coordinates two cooling units as a lead/standby pair, running only one at a time with weekly rotation and automatic failover. A manager enclosure with a status LED and failover shield sits above two unit cards joined by animated command buses; each card carries a gold **LEAD** or gray **STANDBY** badge, spins its fan and glows green while that unit is on, and turns red on that unit's comms alarm. Rotation arrows between the cards show the last changeover, and a pulsing "SWITCHING — Ns" pill rides the bus while the staggered on/off delay counts down. The cards click through to the two managed units.

Dry cooler manager

The dry-cooler manager dashboard is the plant's system view: a status roundel (**NORMAL** / **ALARM** / **OFF**) beside a glycol loop that runs through a cooler-bank glyph with spinning fans, warm return air in and cool supply out, with both loop temperatures shown on the diagram. The supply temperature reddens when a high- or low-supply-temperature alarm is active. A display-only **Controls** section lists the loop **Temperature Setpoint**; clicking it opens that control's own page, where the value is changed.

Pump manager

The pump-manager dashboard is styled like the AC team manager: a **PUMP MANAGER** enclosure with command buses down to one card per real pump, each card showing a mini pump glyph whose impeller spins while that pump runs. A card edges green when running, amber on command/feedback mismatch, red on alarm, and dashed-dim when disabled. An amber **MISMATCH** badge appears when the program's running count disagrees with the pumps' own feedback. Each card clicks through to its pump.

Room and area floor plans

Two dashboards trade the equipment diagram for a scaled plan of a building area, placing live units and readings where they physically sit.

Data center floor plan

This is the most spatial view in the console: a scaled floor plan of the data-center building with a live temperature heatmap. Each cooling unit sits in its real room as a glyph with a spinning fan and live temperature, and casts a colored heat zone — blue where cool, red where hot — clipped to its room so heat never bleeds through a wall; a unit that has lost comms or is disabled goes gray so stale data can't read as a cool spot. Orientation marks (a **BUILDING FRONT** label and a compass) help you place yourself, and clicking a unit or a room label navigates to it. This page refreshes every **6 seconds**.

Admin area floor plan

The Admin Area page is a floor plan of the office suite, drawn as accurate wall outlines with six wall temperature-sensor positions marked as ringed **T** glyphs. It is temporary by design: the sensors are not wired yet, so each shows a blank -- value until the sensor upgrade is finished. The rooftop unit that serves the suite appears as a full-width card above the plan (labeled as being on the roof), with its on/return/supply/outside-air tiles and a return-air tile that reddens on a high- or low-return-air-temperature alarm.

Chapter 7 — Alarms

The **Alarms** area is the console's incident center: it shows you what is wrong right now, what has happened recently, and everything that has ever alarmed on the system. This chapter covers the three alarm lists (Active, Information and History), how to read and filter them, how administrators clear alarms and maintain the archive, and the standalone Alarm View page designed for wall-mounted displays. Configuring *which conditions raise an alarm* is a per-unit task and is covered in Chapter 5 (Unit alarm settings) — this chapter is about the alarm *events* those settings produce.

The three alarm buckets

Click the **Alarms** node in the navigation tree and it expands to three child lists — the DM calls them *buckets*:

Bucket	What it holds	Keeps itself current?	Cleared how
Active	Alarms that are firing *right now*, each with a severity (Critical, Major, Minor or Warning)	Yes — refreshes every 2 seconds	An administrator clears one, several or all; cleared alarms leave the list
Information	Informational events — notices that something happened (a unit was disabled, for example), not fault conditions	Yes — refreshes every 2 seconds	An administrator clears one or several events
History	The permanent archive — every closed alarm occurrence with its start and stop time	No — static archive with a manual Refresh button	An administrator can delete all records, or trim the archive to the last N days

Three points are worth fixing in mind before you work with these lists:

- **Clearing an active alarm does not erase its history.** Active and Information rows are live entries; clearing one removes it from that list only. History is a separate permanent archive and keeps its record of the occurrence.
- **History records cannot be deleted one at a time.** The only maintenance operations on History are **Clear All History** and the keep-last-N-days trim, both Admin-only and both described later in this chapter.
- **History shows the newest 1000 records.** Older records remain in the archive — searches and exports work on what is displayed, but the clear and trim operations act on the *entire* archive, not just the visible 1000 rows.

Every Active and Information alarm carries the same identifying fields: a severity, the time it started, and the **Location**, **Group**, **Unit** and **Type** (the alarm's name on the unit) it came from, plus the alarm **Text**

(the message). History rows carry the same fields with both a **Time Start** and a **Time Stop**.

Severity levels and colors

Active alarms are ranked by severity, and the same four-color code appears everywhere alarms are shown — list rows, badges, the details dialog, the navigation tree and even the browser-tab icon on the standalone Alarm View:

Badge	Severity	Color
CRT	Critical	Red
MAJ	Major	Orange
MIN	Minor	Yellow-gold
WAR	Warning	Teal

In alarm lists, each row shows a small uppercase severity badge, a colored stripe on its left edge and a faint tint of the severity color across the row. **Critical rows pulse** — a slow fade in and out — so they stand out even on a crowded screen. Information events use a blue **INFO** marking instead of the four severities, and the full-page alarm view can additionally show **Cleared** (green) for an alarm that has been resolved.

Severity also drives sorting: when a list is sorted by severity, Critical sorts first, then Major, Minor and Warning — worst news at the top.

The Alarms dashboard

Clicking the **Alarms** node itself opens the alarms dashboard — a simple launching page listing the three buckets.

- The page header shows the title **Alarms** with the same bell icon and color as the tree node, a subtitle counting the containers, and the usual status dot with an "Updated:" time.
- Like other container pages, the dashboard offers **Cards** and **Table** views (see Chapter 3); your choice is remembered by the browser and shared with every other list page in the console. Table view adds a **Description** column.
- Click a card or row to open that bucket. The navigation tree follows along, so **Active**, **Information** or **History** highlights in the tree as its page opens.

The bucket icons on the dashboard use the same state colors as the tree — when active alarms exist, the Alarms bell shows red on the dashboard just as it does in the tree.

Working the Active alarms list

Open **Alarms** → **Active** to see everything that is alarming right now.

The page title **Active** renders in red, and the list refreshes itself **every 2 seconds** — the footer confirms it: "Active Alarms — N alarms | Auto-refresh every 2 seconds". The status dot at the top-right is red whenever at least one alarm is active and green when the list is empty, and the "Updated:" stamp shows the last refresh time. A manual **Refresh** button sits beside them if you want an immediate update.

By default the list sorts by severity, worst first. Every column header is clickable: click once to sort by that column, click again to reverse the direction — the active column shows an orange arrow.

Column	What it shows
Severity	The severity badge (and the alarm's icon, where one is defined)
Time	When the alarm started — click it to open the alarm details dialog
Location	The location the alarm came from, as a link
Group	The group, as a link
Unit	The unit, in bold, as a link
Type	The alarm's name on the unit, as a link to its definition
Text	The alarm message

The **Location**, **Group**, **Unit** and **Type** links are shortcuts into the rest of the console: clicking **CRAC-01** in an alarm row jumps the navigation tree straight to that unit's page, and clicking the **Type** opens the alarm's definition under the unit, where its thresholds live (Chapter 5). If the item an alarm refers to has since been deleted, that cell shows plain text instead of a link.

Note: Administrators also see a checkbox at the left of every row — that column belongs to the clearing workflow described below and does not appear for View users.

Filtering by severity

Below the header sit four legend badges — **CRT Critical**, **MAJ Major**, **MIN Minor**, **WAR Warning**. They double as filters:

- With nothing selected, all severities are shown.
- Click a badge to show only that severity. The selected badge gets a white ring; the others dim. You can select more than one — for example, **CRT** and **MAJ** together to hide the minor noise during an incident.
- Click a selected badge again to deselect it.

Searching

The search bar ("Search alarms...") filters the list as you type, matching any column — severity code, time, location, group, unit, type or message text — without regard to case. An ✕ button appears in the

box to clear the search. The footer counts what you are seeing, for example "3 of 12 alarms", and if nothing matches, the list shows "No alarms match the current filter."

The alarm details dialog

Click any alarm's **Time** link to open its details in a dialog without leaving the list.

The dialog header shows a severity chip in the alarm's color next to the title **Alarm Details**. The body lists the alarm's Time, Location, Group, Unit, Type and Message (fields with no value are simply omitted). The footer offers:

- **Open full page** ↗ — navigates to the alarm's own full-page view (described later in this chapter).
- **Clear Alarm** — Admin only; clears just this one alarm. If the clear fails, an alert reports the reason ("Clear failed:" followed by the message).
- **Close** — dismisses the dialog. Clicking outside the dialog or on the ✕ does the same.

Selecting alarms and clearing them

Admin only: Clearing alarms — singly, selected, or all — requires Administrator access. View users see the list without checkboxes or any clear controls.

To clear several alarms at once:

1. Tick the checkbox on each alarm you want to clear. The checkbox in the header row selects or deselects every visible row at once — "visible" means the rows that pass your current search and severity filter.
2. As soon as one row is selected, an orange **selection action bar** slides in, reading "N alarm(s) selected" with two buttons: a red **Clear Selected** and an ✕ **Cancel** that deselects everything. Selected rows highlight amber (and a selected Critical row stops pulsing so the selection is easy to see).

1. Click **Clear Selected**. A confirmation asks:

"Clear N selected active alarm(s)?"

This cannot be undone."

1. Confirm, and the alarms are removed from the Active list. A brief "Saving..." overlay appears while the console settles, then the navigation tree refreshes so the cleared alarm entries disappear from it too.

Caution: Clearing active alarms cannot be undone — the confirmation says exactly that. The alarms' History records are unaffected, but the live entries are gone. If the underlying condition still exists on the equipment, the alarm can of course fire again.

If a clear fails, an alert reports "Clear selected failed:" with the reason. Two selection behaviors are worth knowing: your selection survives sorting and filtering, but it resets whenever the 2-second refresh delivers fresh data — and to keep a selection from being wiped mid-operation, **auto-refresh pauses while any rows are checked** and resumes when the selection is empty again.

When all is quiet

With zero active alarms, the list is replaced by a green panel with a check icon: "No active alarms — all systems normal." The status dot shows green and the footer reads "0 alarms".

Information events

Open **Alarms** → **Information** for the informational bucket. It works almost exactly like the Active list — same 2-second refresh (with the same pause while rows are selected), same search bar (placeholder "Search events..."), same sortable columns with Location/Group/Unit/Type links, same details dialog from the **Time** link — with these differences:

- **No Severity column and no severity filter** — every entry is informational. Rows carry a blue left edge, and the page's sort arrows and accents are blue rather than orange. The default sort is by **Time**, newest first.
- **The status dot stays green** after a successful load — information events are notices, not faults, so they never turn the indicator red.
- The details dialog shows an **INFO** chip, and its clear button is labeled **Clear Event** (Admin only).
- The selection bar reads "N event(s) selected", and **Clear Selected** confirms with:

"Clear N selected information event(s)?

This cannot be undone."

Caution: As with active alarms, clearing information events is permanent — there is no undo.

If no events are present the page shows "No information events found."; if a search matches nothing, "No events match" followed by your search term in quotes.

Alarm history

Open **Alarms** → **History** for the permanent archive of every alarm that has started and stopped.

History behaves differently from the live lists, on purpose:

- **No auto-refresh.** History is a static archive; the footer reads "Alarm History — N records | No auto-refresh". Click the **Refresh** button in the header when you want the latest records.
- **Newest 1000 records.** The list shows the most recent 1000 entries, newest first, with times in MM/DD/YYYY HH:MM:SS format. Older records stay in the archive even though they are not displayed.

- **Columns: Time Start** (the default sort, newest first), **Time Stop, Location, Group, Unit, Type, Text**. The two time columns sort chronologically — a genuine date sort, not an alphabetical one. Rows carry a neutral gray edge; history rows have no severity.
- **No links, no checkboxes.** History rows do not link into the tree. Instead, **click anywhere on a row** to open a read-only details dialog showing the record's time span ("start → stop") and its Location, Group, Unit, Type and Message. There is no clear button and no full-page link — individual history records cannot be removed.
- **Search** works as on the other lists (placeholder "Search history..."), across all seven columns. No matches shows "No records match" with your term; an empty archive shows "No alarm history found."

Clearing all history

Admin only: The **Clear All History** button and the trim controls described next appear in the toolbar only for Administrator accounts. View users do not see them at all.

Clear All History (trash icon; it turns red as you hover) deletes the entire archive. The confirmation reads:

"Delete ALL alarm history?"

This deletes every record from the ALARMS table and cannot be undone."

Confirm, and every history record — not just the 1000 on screen — is deleted; the list reloads empty. If the operation fails, an alert reports "Clear history failed:" with the reason.

Caution: This is the most destructive action on the alarm pages. Once confirmed, the entire alarm archive is gone and cannot be recovered. If you need a copy first, export the list to PDF or Excel before clearing (see below) — bearing in mind the export contains at most the 1000 displayed records.

Trimming history to the last N days

For routine housekeeping, the trim control keeps recent history and discards the old. Next to **Clear All History** you'll find "Keep last [30] days" with a **Delete Older Records** button.

1. Enter the number of days of history to keep (the box defaults to 30). If the value is blank, not a number, zero or negative, the page responds "Please enter a valid number of days (e.g. 30)." and puts you back in the field.
2. Click **Delete Older Records**. The confirmation reads:

"Delete all alarm history older than N day(s)?"

This cannot be undone."

1. Confirm, and every record older than N days is deleted from the archive — again, across the whole archive, not just the visible rows. Records within the window are untouched.

Caution: Trimming cannot be undone. Choose the day count carefully — the deletion is measured against each record's start time, and anything older is permanently removed.

A failed trim reports "Trim history failed:" with the reason.

Exporting alarm lists

Every alarm list — Active, Information and History — has an export bar with two buttons:

- **PDF** — downloads a formatted PDF report: a title block with "Pythia Technologies" and a timestamp, a note of any filter in effect and the record count, then the table with alternating row shading and a severity color stripe on each row. Long messages wrap, and each page carries a "Page i of n" footer.
- **Excel** — downloads a real `.xlsx` spreadsheet with a few header rows (title, export timestamp, filter and record count) followed by the data.

Both exports contain **exactly what you are looking at**: the current severity filter and search are applied, and a History export contains at most the 1000 loaded records. Filenames identify the list and the moment — for example `Active_Alarms_20260705_0930.pdf`, `Information_Events_...xlsx` or `Alarm_History_...pdf`.

Note: Exports are available to every user, View and Admin alike. If an export button complains that a library is not loaded, contact your administrator — a support file is missing from the DM's web installation.

The full-page alarm view

Each active alarm and information event also has a page of its own. You reach it two ways: click **Open full page** ↗ in the details dialog, or click the alarm's timestamp entry in the navigation tree, where each alarm appears as a leaf under **Active** or **Information**.

The page title states the severity — "Critical Alarm", "Major Alarm", "Information Alarm" and so on — tinted in the severity color, with the alarm's Type as the subtitle. The body is a single detail card with a severity-colored rule down its left edge, listing:

Field	Content
Severity	A badge with the full label — Critical, Major, Minor, Warning, Information or Cleared
Started	The alarm's start time
Location / Group / Unit / Type	Links to the source items in the tree — plain text if the item was deleted, "(none)" if not applicable
Description	The alarm text, or "(none)"
Category	The alarm's category, or "(none)"

The page loads once and does not auto-refresh. The status dot shows red for the four alarm severities and green for informational or cleared alarms.

Administrators see a **Clear** button at the top-right of the card (muted gray until you hover, when it turns red). Clicking it opens a confirmation dialog titled **Clear Alarm**:

"Clear this alarm" followed by the alarm's name in quotes, with the caution line "This action cannot be undone." and **Cancel** / **Clear** buttons.

Confirm, and a brief "Saving..." overlay appears while the console settles; the tree then returns you to the bucket list (Active or Information) with the alarm gone. If the clear fails — for example, another administrator cleared the alarm a moment earlier — the dialog shows the error in red and the button re-enables; an alarm cleared elsewhere shows "Alarm container does not exist" if you try to load or clear it.

The standalone Alarm View for wall displays

The DM includes a separate, self-contained alarm page designed to run full-screen on a wall monitor or an operator's second display: the **Alarm View**. It is a single page in its own browser tab — no navigation tree, no sign-in, no editing controls of any kind — and it is mobile-friendly, so it also works on a phone or tablet on the plant floor. Ask your administrator for its address; it is served directly from the DM alongside the console.

Across the top:

- **Launch DM** — opens the DM Console sign-in page in a separate tab, so an operator watching the wallboard can jump into the full console when something needs attention.
- **Active / History** — the mode toggle. **Active** highlights red and shows the live alarm list; **History** highlights green and shows the archive. Switching modes clears any search or filter and resets the sort.
- **Flash** — a bell-icon pill that controls the new-alarm tab flash (described below). Amber when on; on by default. The setting is remembered by the browser.
- The status dot, "Updated:" time and a manual **Refresh** button, just as on the console pages.

Active mode shows the same live table as the console's Active list — severity badges, pulsing Critical rows, 2-second auto-refresh, search and PDF/Excel export — but strictly read-only: the severity legend is for reference (not clickable filters), rows carry no links or checkboxes, and clicking a row does nothing. With no alarms, the familiar green "No active alarms — all systems normal." panel appears.

History mode shows the archive with manual refresh only, and adds a quick time filter: "Show: **Last 4h** | **Last 8h** | **Last 12h** | **Last 24h** | **All**". **Last 4h** is the default — a wallboard usually cares about the recent past — and the active choice highlights green. Search and exports respect the mode, the time filter and the search term together.

Tab flash and the severity favicon

The Alarm View is built to get attention even when its tab is in the background:

- **The tab icon always reflects the worst active severity.** The browser-tab icon becomes a colored square with a white "!" — red for Critical, orange for Major, yellow for Minor, teal for Warning. With zero active alarms it returns to the normal Pythia icon.
- **When a new alarm arrives** — that is, when the active-alarm count rises — the tab starts flashing: the tab title alternates between "! N Active Alarm(s) - Pythia" and the normal title, and the icon blinks in step, roughly every three-quarters of a second.
- The flash stops on its own when you focus the tab, when active alarms drop to zero, or after 30 seconds. Opening the page fresh never flashes — it takes the current count as its baseline and flashes only on increases from there.
- The **Flash** pill turns the flashing off entirely; the steady severity icon still applies. Switching to History mode also stops any flash and restores the normal icon.

Note: The console's own Active alarms page deliberately does **not** flash the browser tab — inside the console you move between pages constantly, and a changing tab icon would read as noise. The flash-and-favicon experience belongs to the standalone Alarm View alone.

What View users can and cannot do

Everything observational in this chapter is available to both access levels. A View user can open all three buckets and the dashboard, watch the Active and Information lists refresh, filter by severity, search, sort, open the details dialog and the full-page alarm view, export any list to PDF or Excel, and run the standalone Alarm View.

What View users cannot do is remove anything: the row checkboxes, the selection bar, **Clear Selected**, the **Clear Alarm** / **Clear Event** buttons in the dialogs, the **Clear** button on the full-page view, and the History maintenance controls (**Clear All History** and **Delete Older Records**) are all absent for View accounts. Alarms leave the lists only when an administrator clears them — or, in the case of the Active list, the DM itself moves an alarm to History when its condition ends.

Chapter 8 — Automations and Notifications

The DM console can watch your site around the clock, but it only helps if someone finds out when something goes wrong. **Automations** are the rules that carry news out of the console: they send emails and SNMP traps when alarms occur, and they can send scheduled "system is alive" heartbeat emails so that silence never gets mistaken for health. This chapter covers creating automations, managing their email recipients, SNMP trap destinations and time exceptions, and setting up heartbeat schedules. Everything here is configured by Administrators; View users can look but not change.

What automations do

An automation is a named rule under the **Automations** branch of the navigation tree. There are two types, chosen when the automation is created:

- **Alarm** — fires when alarms occur. Each Alarm automation carries three fixed action lists:
- **Email** — the recipients who receive an alarm email.
- **SNMP Trap** — the network management stations that receive an SNMP trap.
- **Time** — *time exceptions*: day-of-week entries, each holding one or more start/stop time ranges, that constrain *when* the automation's notifications apply (for example, only 08:00–17:00 on weekdays). With no time entries, the automation applies at all times.
- **Heartbeat** — sends periodic proof-of-life emails on a schedule rather than in response to alarms. Heartbeats are covered in their own section later in this chapter.

Every automation can be enabled or disabled as a whole. A disabled automation never fires, and the console shows it dimmed everywhere so you can tell at a glance that it is switched off.

You can create as many automations as you need. A common arrangement is one broad automation (say, **ALL ALARMS** emailing the operations inbox), a narrower one for after-hours escalation (**CRITICAL ONLY** with a pager gateway recipient and weekday time exceptions), and a **HEARTBEAT** automation feeding an upstream monitoring system.

The Automations list

Click the **Automations** node in the navigation tree. The page is headed **Automations**, with a subtitle counting what it found — "3 Automations", for example. On a new system the list is empty:

Like other container pages, the list offers **Cards** and **Table** views (your choice is remembered — see Chapter 3). In Cards view each automation shows its name, its type (**Alarm** or **Heartbeat**) and its description if one is set; disabled automations render at half opacity. Clicking a card or row opens that automation.

The Table view adds something the cards do not have: an inline **Enabled** toggle per automation.

Note: This page loads once when you open it and does not auto-refresh — automations change rarely. There is no search bar here either; the lists *inside* an automation have one.

Enabling and disabling an automation

Admin only: View users see the Enabled column as plain "✓ Yes" / "— No" text.

In Table view, click the slider in the **Enabled** column. It flips immediately, dims briefly while the change saves, then re-enables. If the save fails, the slider flips back and a message explains why (or simply "Enable toggle failed.").

This is the way to silence an automation without deleting it — the automation and all its recipients, destinations and time settings stay configured, ready to switch back on. The Cards view has no toggle; use the Table view, or the automation's own Properties dialog.

Adding an automation

1. Click **Add** in the section toolbar. The **Add Automation** dialog opens.
1. Enter an **Automation Name** (required, up to 100 characters — for example `ALL ALARMS`).
2. Choose the **Type: Alarm** (the default) or **Heartbeat**. This choice is fixed at creation — the type determines which action lists the automation is built with.
3. Optionally enter a **Description** (up to 200 characters), and click **Add**.

If the name is missing you are told "Automation name is required."; a name or description containing the pipe character is refused with "Name and description cannot contain the pipe character (|)." A name already in use answers "An automation with that name already exists".

On success the console builds the automation's entire fixed structure in one step — for an Alarm automation the **Email**, **SNMP Trap** and **Time** lists; for a Heartbeat, its **Interval** and **Period** schedules — shows the brief "Saving..." overlay, and then navigates straight to the new automation so you can start adding recipients.

The From address

Outgoing automation emails are stamped with a reply-to **From** address, set once for all automations:

1. On the Automations list page, click **Properties** in the toolbar. The **Automations Properties** dialog opens.
2. Enter the **From** address — for example `alerts@example.com` — and click **Save**.

A non-empty value must be a valid email address ("From address must be a valid email (e.g. user@example.com)."). Leaving the field empty is allowed; the system then falls back to its standard SMTP reply address. The same dialog also carries optional **SMTP Relay** and **SMTP Port** fields that override the local mail relay's smart-host and port; leave them blank to use the relay's defaults (port 25,

no smart host). These per-automation fields override the shared mail relay for this automation's messages only; the relay itself — its queue, delivery status, and allow list — is configured on the **SMTP Relay** page (see **SMTP Relay** under System Settings).

Inside an Alarm automation

Click an Alarm automation on the list (or its tree node). The page is headed with the automation's name and the prompt "**Select an action type below**", and shows the three fixed action lists as cards:

Action	What it holds
Email	List of Email Recipients
SNMP Trap	List of SNMP Trap Destinations
Time	List of Time Exceptions

These three are always present and cannot be added to or removed — you click into each one to manage its contents. The next three sections cover them in turn.

Automation properties, rename and delete

The **Properties** button on this page opens the **Automation Properties** dialog for the automation itself: the name with the rename pencil at the top (Chapter 3 covers the pencil flow), the **Description** and other property fields below, and the red **Delete** button in the bottom-left.

Renaming follows the standard name rules (1–100 characters, none of \ & % # ? " , , no duplicate). When the rename succeeds, everything beneath the automation — its Email, SNMP Trap and Time lists and all their contents — keeps working under the new name; the tree, the page title and the address all update together.

Deleting an automation is always a cascade — an automation always has contents. Clicking **Delete** opens the **Delete Automation** confirmation:

> Are you sure you want to delete the automation "ALL ALARMS" and everything inside it (all recipients, trap destinations, and time exceptions)?

> This action cannot be undone.

Caution: Confirming removes the automation *and* every recipient, trap destination and time exception configured under it, permanently. If you only want to stop it firing, disable it instead.

After a successful delete the console shows the "Saving..." overlay briefly and returns you to the Automations list.

Email recipients

Click the **Email** card inside an Alarm automation. The **Email Recipients** page lists everyone who receives this automation's alarm emails, with a search box ("Search recipients...") and a manual **Refresh** button.

Each recipient is identified by its email address and carries three flags plus an optional description:

- **Enabled** — whether the recipient is active. Disabled recipients show a gray dot and render dimmed.
- **Short Format** — sends this recipient a condensed, text-message-style body instead of the full email. Shown as a blue **Text** badge on the card. Use this for recipients that are really SMS or pager email gateways, where a long body would be truncated or split.
- **Test** — marks a test recipient, who only receives test sends. Shown as an orange **Test** badge.

In Table view the columns are **Email**, **Enabled**, **Short Format**, **Test** and **Description**; every column header sorts the list, and flag columns show "✓ Yes" or "— No".

A brand-new automation starts with no recipients:

Adding a recipient

1. Click **Add**. The **Add Email Recipient** dialog opens.
2. Enter the **Email Address** (required) — for example `oncall@example.com`.
3. Optionally add a **Description** (such as "On-Call Pager") and set the checkboxes: **Enabled** is pre-checked; tick **Short Format** for an SMS/pager gateway; tick **Test** for a test-only recipient.
4. Click **Add**.

The address is checked before anything is sent: an empty field shows "Email address is required.", and a malformed one shows "Must be a valid email (e.g. user@example.com)."

An address already on this automation's list answers "A recipient with that email already exists". On success the console navigates to the new recipient's own page.

Editing or deleting a recipient

Click a recipient on the list to open its own properties page — the standard leaf property page from Chapter 3. From its **Properties** dialog you can change the Enabled, Short Format and Test flags and the description, or click **Delete** to remove the recipient. The email address itself is the recipient's name; renaming it follows the pencil flow, and the new name must be a valid email address.

SNMP trap destinations

Click the **SNMP Trap** card inside an Alarm automation. The **SNMP Trap** page lists the network management stations that receive a trap when this automation fires — the page works exactly like the recipients page, with cards or a table, search ("Search destinations...") and a manual Refresh.

Each destination is identified by its IPv4 address and carries:

- **Enabled** — whether the destination is active; disabled destinations render dimmed.
- **Port** — the UDP port the trap is sent to: **162** (the SNMP standard) or **8162** (the common alternate for receivers that cannot bind the privileged port).
- **Persistent** — retry the trap until the receiver acknowledges it.
- **SNMP v2** — send an SNMPv2c trap instead of v1. Set this to match what your management station expects.
- **Test** — marks a test destination.
- **Description** — optional, e.g. "NMS Receiver".

In Table view the columns are **IP Address**, **Enabled**, **Port**, **Persistent**, **SNMP v2**, **Test** and **Description**, all sortable.

Adding a trap destination

1. Click **Add**. The **Add SNMP Trap Destination** dialog opens.
1. Enter the **IP Address** (required) — for example `192.0.2.10`.
2. Pick the **Port** (162 by default), optionally add a **Description**, and set the checkboxes — **Enabled** is pre-checked.
3. Click **Add**.

An empty address shows "IP address is required."; anything that isn't a dotted-quad shows "Must be a valid IPv4 address (e.g. 192.168.1.10)." A duplicate answers "A destination with that IP already exists". The Port dropdown only offers valid values, so it never needs correcting.

Like recipients, individual destinations are edited or deleted from their own properties page — click the destination on the list to open it.

Time exceptions

By default an automation applies at all times. The **Time** list narrows that: you add day-of-week entries, and inside each day one or more start/stop time ranges. When any time entries exist, the automation's notifications apply only inside those windows — for example, a paging automation limited to business hours, or an after-hours automation covering evenings and weekends.

The day list

Click the **Time** card inside an Alarm automation. The **Time** page lists the configured days, named in the console's stored convention `1 - SUNDAY` through `7 - SATURDAY` (the digit keeps them in week order). With none configured the page reads "No time exceptions configured."

To add a day, click **Add**, pick the day from the dropdown in the **Add Day of Week** dialog, and click **Add**. Each day can appear only once per automation — adding a day that already exists is refused with:

> "That day is already added to this automation"

On success the console navigates into the new day so you can add its time ranges.

Time ranges within a day

Click a day to open its time-range page, headed with the day's name. Each range is shown as `start → stop` in 24-hour time, with an optional description:

To add a range:

1. Click **Add**. The **Add Time Range** dialog opens.
2. Pick the **Start Time** and **Stop Time** — hourly steps, start from 00:00, stop up to 24:00 — and optionally a **Description** ("Morning shift").
3. Click **Add**.

The dialog enforces the obvious ordering rule — "Stop time must be after start time." — and, on the server side, that ranges within one day never overlap. Attempting to add a range that collides with an existing one is refused with a message naming the conflict:

> "This time range overlaps the existing range 08:00 - 12:00"

To cover a split shift, add two non-overlapping ranges as in the figure above. Two ranges also cannot share a start time ("A range with that start time already exists"). Click a range on the list to open its properties page, where the stop time and description can be edited and the range deleted; the start time is the range's name and renames via the pencil.

Renaming or deleting a day

The day page's **Properties** button opens a dialog titled **Properties:** followed by the day (for example **Properties: 2 - MONDAY**) with a single day dropdown and the red **Delete** button. Choosing a different day and clicking **Save** moves the entry — with all its time ranges — to that day; picking a day that already exists on this automation is refused with "That name already exists".

Clicking **Delete** opens the **Delete Day** confirmation:

- > Delete 2 - MONDAY and any time ranges configured under it?
- > This action cannot be undone.

Confirming removes the day and every range inside it, then returns you to the day list.

Heartbeat automations

Alarm automations only speak up when something goes wrong — so if the DM server itself dies, loses its network link, or its mail path breaks, you hear nothing, and silence looks exactly like health. A **Heartbeat** automation closes that gap: the console sends "system is alive" emails on a fixed schedule, and the recipients — a person, a ticketing inbox, or an upstream monitoring system — expect them on time. **A missing heartbeat is itself the warning sign:** it tells you the DM host, its automation engine, or its email path has stopped working, end to end.

Create a Heartbeat automation from the same **Add Automation** dialog, choosing **Heartbeat** as the **Type**. Clicking it on the Automations list opens its landing page, which offers its two fixed schedule types:

Schedule type	What it does	Typical audience
Interval	Repeating emails around the clock — every 5 minutes, 15 minutes, or 1 hour	Monitoring systems that watch for a missed beat
Period	One email per day at a chosen top-of-hour time (e.g. 05:00 daily)	People who want a daily "all is well" note without inbox flooding

A single heartbeat automation can use both at once — for example a 5-minute interval feeding a monitoring system, plus an 06:00 daily period emailing the facilities manager.

Intervals — three fixed slots

Click the **Interval** card. The Interval page always lists the same three entries, created automatically with the automation:

- **01 (5-Minutes)**
- **02 (15-Minutes)**
- **03 (1-Hour)**

These slots are fixed — they cannot be added to, renamed or deleted, by anyone. You simply click into whichever cadence you want and add recipients to it (below). A slot with no recipients sends nothing.

Interval emails fire on the even divisions of the clock: the 5-minute slot at :00, :05, :10 and so on; the 15-minute slot at :00, :15, :30 and :45; the 1-hour slot on the hour.

Periods — daily emails at a chosen hour

Click the **Period** card. Unlike intervals, periods are entirely yours to manage: each period is a daily send time, named by that time (00:00 through 23:00), and you add as many as you need.

To add one, click **Add**, pick the hour from the **Period** dropdown in the **Add Period** dialog, and click **Add**. Each time can appear only once — a duplicate is refused with "That period is already added to this automation". A period's email goes out each day when the clock reaches that hour.

To change or remove a period, open the period (click it on the list) and use the **Properties** button on its recipients page — this button appears only for periods, since interval slots are fixed. The dialog, titled **Properties:** followed by the time (for example **Properties: 05:00**), has a single time dropdown and the red **Delete** button. Picking a new time and saving moves the period — recipients and all; a time already in use is refused with "That name already exists". Clicking **Delete** opens the **Delete Period** confirmation:

> Delete 05:00 and any email recipients configured under it?

> This action cannot be undone.

Recipients per schedule slot

Every interval slot and every period has its own recipient list — click the slot to open it. The page is headed with the slot's name (for example **01 (5-Minutes)** or **05:00**) and works exactly like the alarm-automation recipients page: cards or table, search, manual Refresh, and the same **Add Email Recipient** dialog with the **Email Address**, **Description**, **Enabled**, **Short Format** and **Test** fields and the same validation messages.

Because each slot has its own list, the same address can receive, say, the hourly interval but not the 5-minute one — add it only where you want it. Each entry carries an **Enabled**, a **Short Format** (shown as **Text**), and a **Test** flag. **Short Format** is honored on heartbeats: a recipient marked **Text** receives the compact message described below, and everyone else receives the full one. To change a recipient's flags or remove it, click the entry to open its properties page, the same way as for alarm recipients.

Caution: On heartbeat schedules, the **Enabled** checkbox on a recipient is **not** enforced by the sender — a recipient still receives heartbeat emails even when Enabled is switched off. To actually stop heartbeats going to someone, **remove them from the slot's recipient list** rather than just unchecking Enabled.

What the heartbeat emails look like

A normal heartbeat email is short and self-describing:

- **Subject:** DM Console Heartbeat - <Site> - <YYYYMMDDHHMM>.<Interval|Period> — your site name (if one is configured in the system settings), a timestamp, and which schedule type fired. For example: DM Console Heartbeat - Sunbury Office - 202607050900.Interval.
- **Body:** a matching single line — Heartbeat message from DM Console - <Site> - <timestamp>.<type>.

Recipients flagged **Short Format** instead receive the compact variant: the subject is just `Heartbeat`, and the full subject line above becomes the body — sized for SMS and pager email gateways.

The From address on heartbeat emails is the **From** value set on the Automations root (see "The From address" earlier); if that is empty, the system's standard SMTP reply address is used.

Note: If your recipients monitor for missed heartbeats, remember that disabling the heartbeat automation (or deleting a period) stops the emails by design — tell the people or systems watching for them first, or the silence will be read as an outage.

Renaming or deleting a heartbeat automation

The heartbeat landing page's **Properties** button works like the Alarm automation's: rename with the pencil, edit the **Description** (the **Type** field is shown read-only), and the red **Delete** button. Deleting is

again a cascade, with its own wording in the **Delete Automation** confirmation:

> Are you sure you want to delete the heartbeat automation "HEARTBEAT" and everything inside it (all interval and period schedules)?

> This action cannot be undone.

Caution: Confirming removes every interval and period schedule and all their recipients, permanently.

What View users see

View users can browse everything in this chapter read-only: the Automations list (with Enabled shown as plain "✓ Yes" / "— No" text instead of toggles), each automation's action lists, recipients, trap destinations, time exceptions and heartbeat schedules, with search and sorting working normally. No **Add**, **Properties** or **Delete** buttons appear anywhere in the Automations branch, and the interval pages — read-only even for Administrators — look identical for both roles.

Chapter 9 — Trends

Trends are the DM's historical record. While the **Live Data** pages show what a point reads *right now*, a trend stores samples of that point over time so you can chart yesterday's temperature swing, last month's UPS load, or a season of daily readings. Operators use the Trends pages to review and export history; administrators additionally choose which data points get trended and can delete a trend's history.

How trending works

The DM collects trend history in **five fixed time buckets**, one per sampling interval. The buckets always exist — you don't create or delete them — and each holds the trends sampled at that rate:

Bucket	Samples taken	Chart window shown in
01 (5-MINUTES)	every 5 minutes	Hours
02 (15-MINUTES)	every 15 minutes	Hours
03 (30-MINUTES)	every 30 minutes	Days
04 (1-HOUR)	every hour	Days
05 (1-DAY)	every day	Days

A few rules govern the model:

- **One bucket per point.** A data point is trended into at most one bucket at a time. Assigning it to a bucket starts collection at that rate; setting it to **Off** stops collection.
- **Trends create themselves.** Once you assign a point to a bucket, the trend service creates the trend entry under that bucket automatically — it appears within a short sync pass. You never add a trend by hand.
- **History survives renames.** Each trend carries a permanent identity (a GUID, visible in the trend's **Properties**), and the collected values are stored in the DM's database under that identity. Renaming the location, group, or unit that the point belongs to does not break or discard its history.
- **Collection never sleeps.** Samples are recorded around the clock whether or not anyone has a trend page open. You can walk away for a month and come back to a full chart.
- **Trend names describe their source.** A trend is named after the full path of its data point — for example `_PTC_Locations_Sunbury_Office_Environmentals_Office_TH1_Data_Temperature` — so you can always tell what it records from the name alone.

The Trends dashboard

Click **Trends** in the navigation tree to open the Trends dashboard. The header shows the page title with a subtitle counting the buckets, the status dot and "Updated" time on the right, and the blue **Configure Points** button (which manages trending system-wide — see the next section).

The **Trend Time Buckets** section lists the five buckets, in either Cards or Table view (the toggle works the same as elsewhere in the console, and your choice is remembered):

- **Cards view** — one card per bucket with its name (e.g. **01 (5-MINUTES)**), a description such as "5 Minute Trends", and a poll-rate badge showing the sampling rate in seconds (e.g. `300s`).
- **Table view** — columns **Bucket**, **Description**, and **Poll Rate (s)**.

Click any bucket card or row to open that bucket and browse the trends inside it. This page loads once and does not auto-refresh; re-click the **Trends** node in the tree to reload it.

Choosing which points to trend

Click **Configure Points** on the Trends dashboard to open the **Configure Trended Points** dialog — a single board that covers every trendable data point in the system. The subtitle summarizes it at a glance: "N trendable points · M trended", with "(showing K)" appended whenever a filter is narrowing the list.

Admin only: assigning and clearing trend buckets requires Administrator access. View users can open the dialog and browse or filter it, but the bucket dropdowns are disabled and the selection checkboxes and bulk-apply bar are hidden.

The table groups points by **Location > Group > Unit**, with a collapsible header row per group showing "N points · M trended" (and "K selected" when you have rows checked in it). Each point row shows:

- **Point** — the point name in bold, with its description underneath when the two differ.
- **Value** — the point's current live value, so you can confirm you have the right point before trending it.
- **Trend bucket** — a dropdown offering **Off** plus the five bucket names. The dropdown gets a blue border when a bucket is set, making trended points easy to spot.

Assigning a single point

1. Type in the **Search** box ("Filter by point, unit, group, or location...") to narrow the list. The filter is live and matches point names, descriptions, units, groups, and locations.
2. In the point's row, pick a bucket from the **Trend bucket** dropdown.
3. The change saves immediately — a green "✓ Saved" flashes next to the dropdown. There is no separate Save button.

The trend service then creates the trend entry in the chosen bucket in the background. To stop trending a point, set its dropdown back to **Off**.

Note: you can also set the same **Trend** property from the data point's own properties dialog (Chapter 5); it offers the same five buckets. The Configure Points dialog is simply the system-wide view of that property.

Assigning many points at once

For bulk changes, use the checkbox column:

1. Check the rows you want to change, or use the select-all checkbox in the table header. Select-all applies to **every point matching the current filter**, including points inside collapsed groups — so filter first if you don't mean everything.
2. In the dialog footer, the "N selected" counter tracks your selection. Pick a bucket (or **Off**) from the footer dropdown.
3. Click **Apply**.

Two situations ask you to confirm before proceeding:

- Turning trending off always confirms: "Turn OFF trending for N points?"
- Setting a bucket on **more than 20 points** confirms: "Set Trend "<bucket>" on N points?"

While the changes save, the footer counts "Saving i/N..." and the dialog cannot be closed. On success it reports "Saved N points." If some rows fail, the message reads "Saved X, failed Y (...). Failed rows stay selected." — the failed rows remain checked so you can retry them, while the successful rows are unchecked.

Reviewing what's already trended

Check the **Trending only** box in the dialog toolbar to hide everything except points currently assigned to a bucket. This is the quickest way to audit or prune trending across the whole system — combine it with the search box to review one location at a time.

Browsing a bucket's trends

Open a bucket from the Trends dashboard (or from the navigation tree) to list the trends collected at that rate. The header reads **Trends (<bucket>)** with a count of the trends inside, and the list offers the usual Cards/Table toggle:

- **Cards view** — a card per trend showing its name and a badge with its current viewing window (e.g. `120 Hours`).
- **Table view** — columns **Trend** and **Window**.

Click a trend to open its chart. If nothing has been assigned to this bucket yet, the page shows "No trends in this bucket." There is nothing to add or delete on this page — trends are created through **Configure Points** and deleted from the trend's own page.

Reading the trend chart

Opening a trend shows its collected history as an interactive chart on a dark canvas. The subtitle summarizes what you are looking at — for example "Trend in 01 (5-MINUTES) — 120 Hours — N buckets (with min/max band)".

What the chart shows:

- **The line is the average.** For long spans, the DM consolidates samples so the chart stays responsive; each plotted point is the average of the samples it consolidates.
- **The min/max band.** When consolidation combines multiple samples per plotted point, a translucent band behind the line shows the true minimum and maximum over that interval — so a brief spike is never hidden by averaging. With few enough samples that nothing is consolidated, the band does not appear.
- **The latest dot** marks the most recent sample (not shown in Bar mode).
- **Axis labels** adapt to the span: times only for short windows, date + time for medium windows, dates for long ones. The Y-axis is titled with the source data point's description. Times are the DM's local time.
- A brand-new trend shows "No trend data in this window" until the service has collected its first samples — check back after the bucket's sampling interval has passed.

Hover anywhere on the chart to get a crosshair and a tooltip with the full timestamp (`MM/DD/YYYY HH:MM:SS`) and the exact value.

The toolbar

Left to right across the top of the chart:

1. **Refresh** — re-fetches the history. The trend page does not auto-refresh.
2. **Chart** — Line, Area, Bar, or Step rendering of the same data.
3. **Window** — how far back to display: **24 / 48 / 72 / 120 / 168 Hours** for the 5- and 15-minute buckets, or **30 / 60 / 90 / 120 / 180 Days** for the 30-minute, 1-hour, and 1-day buckets. The default is 120.
4. **Y-Min / Y-Max** — override the vertical scale. Each accepts a number or "auto" (the placeholder). If you enter a minimum that isn't below the maximum, both silently revert to auto and the inputs clear.
5. **Theme** — line color: Blue (default), Cyan, Green, Amber, Magenta, or Mono.
6. **Export** — CSV / Excel / PNG / PDF (next section).
7. **Reset Zoom** (appears only while zoomed) and **Properties**.

Note: the **Window** is a *viewing* preference, not a retention setting. Choosing 24 Hours doesn't discard older data — it only changes how much the chart asks for. Your Window, chart type, Y bounds, and theme choices are remembered per trend in your own browser; they are not shared with other users or other computers.

Zooming in

To examine a detail, click and drag horizontally across the chart. A selection band paints as you drag; releasing zooms the time axis to that range. Repeat to zoom further. The yellow-bordered **Reset Zoom** button appears in the toolbar while zoomed — click it to return to the full window. Changing the **Window** also clears any zoom. Zoom is temporary and is not remembered between visits.

Exporting a trend

The four **Export** buttons save the trend as a file named after the trend and its bucket:

- **CSV** — a `Timestamp,Value` file with one row per plotted point, full timestamps included. Suitable for spreadsheets and further analysis.
- **Excel** — the same data in a file that opens directly in Excel.
- **PNG** — an image of the current chart, rendered at high resolution on the dark background.
- **PDF** — a landscape page with the trend name, the bucket and window, an "Exported" date line, and the chart image.

Note: data exports contain the **average series only** — the plotted line. The min/max band values are not included in CSV or Excel output. To capture the band, export PNG or PDF, which reproduce the chart as displayed.

Exports respect your current view: the selected window, and for image exports the current chart type, zoom, and theme. If the chart hasn't loaded data yet, the buttons do nothing.

Deleting a trend

Deleting a trend removes its collected history *and* stops the point from being trended. Do this from the trend's own page:

1. Click **Properties** in the toolbar. The dialog shows the trend's name and its GUID (the identity under which its history is stored).
2. Click the red **Delete Trend** button.
3. A confirmation dialog opens: "Delete \<trend name\>?" with the caution text: "All historical data values collected for this trend will be permanently deleted. The data point will also stop being trended (you can re-enable it from the data point properties). This action cannot be undone."

4. Click **Delete Trend** to confirm (the button reads "Deleting..." while it works). The tree refreshes and you are returned to the bucket page.

Caution: this permanently erases every stored value for the trend. Re-enabling trending on the same point later (from **Configure Points** or the point's properties) starts a brand-new trend with a fresh identity — the old history cannot be recovered. If you only want to *pause* collection while keeping the history viewable, set the point's Trend bucket to **Off** in Configure Points instead of deleting the trend.

Admin only: deleting a trend requires Administrator access; the DM rejects the request otherwise.

What View users see

View users get the full read side of Trends: the dashboard, bucket lists, every chart control (window, zoom, chart type, theme, Y bounds), and all four exports work identically. In the **Configure Trended Points** dialog they can browse, search, and use **Trending only**, but the bucket dropdowns are disabled and there is no checkbox column or bulk-apply bar. Deleting a trend is refused by the server for View users.

Chapter 11 — Users and Accounts

Every person who uses the DM console signs in with their own account, and every account is one of two kinds: **Admin** or **View**. This chapter covers everything to do with those accounts — the **Users** page where administrators manage them, adding and removing users, changing names and passwords, enabling and disabling sign-in, and the special rules that protect the built-in **Admin** account. It also introduces the **Logons** log, the audit trail that records every sign-in and account change. Signing in itself is covered in Chapter 2; this chapter is about managing the accounts people sign in with.

The two access levels

Each account's **Access** property is either **Admin** or **View** — there are no other levels:

	Admin	View
See monitoring pages, alarms, trends	Yes	Yes
Add, rename, configure, delete anything	Yes	No
Send control commands	Yes	No
Manage other user accounts	Yes	No
Edit their own account (name, password, sign-in email)	Yes	Yes

A View user sees the same console as an administrator, minus the editing controls: no **+ Add** buttons, no **Properties** buttons on system pages, no toggles. The role badge at the right end of the status bar — **Admin** in blue, **View** in gray — always tells you which kind of session you are in.

The one place a View user *can* edit is their own account: they can change their first and last name, their password, and their sign-in email. Everything else about their account — whether it is enabled, and its access level — is set by an administrator.

The Users page

Click the **Users** node in the navigation tree to open the user list. This is the home page for account management: every account on the DM, one per row.

The header shows the title **Users**, a count ("4 users"), and a status dot. The table has five columns:

- **ID** — the account's sign-in email address, shown in blue with a user icon.
- **Enable** — a toggle switch, green when the account can sign in.
- **First Name** and **Last Name** — the person's name.
- **Access** — **Admin** or **View**.

Click anywhere on a row to open that account's page (see "The account page" below). The footer reminds you: "Click a row to view or edit an account. Toggle Enable inline."

Admin only: The full user list, the **+ Add User** button, and the Enable toggles are administrator features. A View user who opens the Users page sees **only their own row** — other accounts are never even sent to their browser — with no Add button and their Enable toggle grayed out. Their footer reads: "You can edit your own account. Other accounts are managed by an administrator." Clicking their row opens their own account page for self-service.

The inline Enable toggle

Administrators can enable or disable an account directly from the list — no dialog needed. Click the **Enable** switch: it flips immediately and dims for a moment while the change is saved. If the server refuses the change, the switch flips back and a red banner above the table explains why.

The toggle on the **Admin** account's row is permanently grayed out — the recovery account can never be switched off (see "The Admin account" below).

For what disabling actually does, see "Disabling an account" later in this chapter.

Adding a user

Admin only: Only administrators can create accounts.

1. On the **Users** page, click **+ Add User**. The **+ Add to USERS** dialog opens.
2. Fill in the fields, top to bottom:
 - **User Identifier (Email Address)** — the address the person will type at sign-in. It must be a valid email address; this is the account's identity everywhere in the console. Capitalization does not matter — sign-in names are case-insensitive.
 - **Enable User Account** — checked by default; clear it to create the account disabled.
 - **User's First Name** and **User's Last Name** — both required.
 - **User's Password** and **User's Password Confirmation** — the person's initial password, typed twice.
 - **Access Level** — **Admin** or **View**.

1. Click **+ Add User**. The button stays disabled until the form is valid.

The form validates as you type. If the two password fields differ you will see "Passwords do not match."; a malformed address shows "User Identifier must be a valid email address."; and no field may contain the `|` character ("Fields cannot contain the | character."). The server applies the same checks and a few more — for example "A user with that ID already exists" if the address is taken.

When the account is created, the dialog closes and the new row appears in the list. Give the person their email address and initial password, and they can sign in immediately (Chapter 2).

Note: The initial password you type here is sent to the gateway over the secure (encrypted) connection and stored only as a protected hash — it is never kept in readable form. New users can change it to one of their own at any time through **Change Password** (described below).

Note: There are currently no password complexity or length rules — any non-empty password is accepted. Choose strong initial passwords, and encourage new users to change their password after their first sign-in.

The account page

Clicking a row on the Users page opens that account's own page. The header shows the account's sign-in email (for example **ID: jim@example.com**), a property count, and a status dot. Below it, a three-column table — **Property | Value | Description** — lists the account's fields: the ID, Enable, First Name, Last Name, and Access. Boolean values read Yes or No.

This is the same properties-page pattern introduced in Chapter 3: the table is read-only, and all editing happens through the **Properties** button in the toolbar. The page does not auto-refresh, but the Properties dialog re-reads the account every time it opens, so it never shows stale values.

The same page serves two audiences:

- **An administrator** viewing any account gets the full page and can edit every field, rename the account, reset its password, or delete it.
- **Any user viewing their own account** — including View users — gets the **Properties** button too, but scoped to self-service: they can change their name, their sign-in email, and their password. They cannot see or change their own **Access** or **Enable** fields, and there is no Delete — you cannot demote, disable, or delete yourself.

The Properties dialog

Click **Properties** to open the editing dialog.

- The first row shows the **ID** — the sign-in email — with a small **pencil** icon for renaming it (next section).
- Below it are the editable fields: **Enable**, **First Name**, **Last Name**, and **Access**. When you are editing your own account, the Enable and Access rows are hidden; on the **Admin** account they appear but are read-only.
- There is deliberately **no password field** here. Passwords are only ever set through the **Change Password** dialog, described below.
- The footer holds **Delete** (red, administrators only) and **Change Password** on the left, and **Cancel** / **Save** on the right.

Changing a user's name

To correct or update the person's name:

1. Open the account's page and click **Properties**.
2. Edit **First Name** and/or **Last Name**. Neither may contain a backslash, and no value may contain the `|` character.
3. Click **Save**. The dialog closes and the page reloads with the new values.

Any user can do this on their own account; administrators can do it on anyone's.

Renaming a sign-in email

The account's ID is what the person types at sign-in, so renaming it is more consequential than changing a display name — the console treats it accordingly.

1. In the Properties dialog, click the **pencil** next to the ID.
2. The row becomes an edit field. Type the new email address, then press **Enter** or click the green ✓ to confirm (Esc or the red × cancels).
3. The new address is validated as an email — a bad value shows "Must be a valid email (e.g. user@example.com)", and addresses may not contain the characters `\ & % # ? " , .`. If another account already uses the address you will see "That name already exists".
4. On success the name updates everywhere at once — the page header, the navigation tree, and the Users list.

From that moment the account signs in with the new address (and the old one no longer works). The password is unchanged.

Note: If you rename **your own** sign-in email, sign out and sign back in with the new address straight away — your session is tied to your sign-in identity, and it will not outlive the old name for long.

The pencil never appears on the **Admin** account — its name is permanent (see "The Admin account" below).

Passwords

Password changes always go through the **Change Password** button in the Properties dialog footer. The dialog that opens takes one of two forms, chosen automatically.

Changing your own password

When the account is your own — whatever your access level — the dialog is titled **Change Password** and asks for three things: your **Current password**, your **New password**, and **Confirm new password**.

1. Type all three fields. The dialog checks them as you go: "New password and confirmation do not match." if the last two differ, and "New password must differ from current password." if you re-entered the old one.
2. Click **Change Password**. The button reads **Working...** for a second or so — the same deliberate pause as at sign-in, while your browser converts the passwords into cryptographic proofs so they never cross the network in readable form.
3. On success the dialog confirms: "Password changed. Use your new password the next time you sign in." — and closes by itself a moment later. Your current session continues; you are not signed out.

If the current password is wrong, the field clears and refocuses with "Current password is incorrect." The same message can appear if you left the dialog sitting for more than a minute before submitting — just try again.

Resetting another user's password

Admin only: Resetting someone else's password requires Administrator access.

When an administrator opens Change Password on **someone else's** account, the dialog is titled **Reset Password for** followed by the account's email, and asks only for **New** and **Confirm** — there is no current-password field, because the administrator's authority stands in for it. This is how you recover an account whose owner has forgotten their password.

Click **Set Password**; on success the dialog confirms: "Password set. The user can sign in with the new password." Tell the person their new password — and encourage them to change it to one of their own at their next sign-in.

An administrator changing their **own** password always gets the three-field self dialog — even administrators must prove they know their current password.

Disabling an account

Disabling is the reversible way to shut off access — for a contractor between site visits, or an employee on leave. Flip the account's **Enable** toggle on the Users page (or clear the **Enable** field in its Properties dialog).

A disabled account cannot sign in. The person sees the same generic "Invalid credentials" message as a wrong password — deliberately, the sign-in page never reveals whether an account exists or has been disabled. If the person is signed in at the moment you disable them, their current screen does not go blank instantly, but their very next action fails; re-enable the account and they can sign in again, with nothing else changed.

You cannot disable your own account, and the **Admin** account can never be disabled.

Deleting an account

Admin only: Only administrators can delete accounts.

When an account should be removed permanently:

1. Open the account's page and click **Properties**.
2. Click the red **Delete** button in the dialog footer.
3. A confirmation dialog asks "Delete **jim@example.com**?" and warns: **This action cannot be undone**.
4. Click **Delete** to confirm (or **Cancel** to keep the account).

The account disappears from the tree and you are returned to the Users list.

Caution: Deleting an account is permanent — there is no undo and no recycle bin. If you only need to stop someone signing in, disable the account instead; you can re-enable it later with everything intact.

You cannot delete your own account — the Delete button never appears when you are viewing yourself — and the **Admin** account can never be deleted.

The Admin account

Every DM has a permanent built-in account whose sign-in name is simply `Admin`. It is the **recovery account**: the one identity that is guaranteed to exist, be enabled, and hold Administrator access no matter what else happens to the user list. If every other administrator account were deleted, disabled, or demoted — by accident or otherwise — `Admin` is how you get back in and repair things. To keep that guarantee, the console enforces special rules on it, in both the pages and the server:

Action	Regular account	Admin account
Rename the sign-in ID	Yes — administrator or the user themselves	Never — the rename pencil does not appear
Disable	Administrator can	Never — the Enable toggle is grayed out
Change Access level	Administrator can	Never — pinned to Admin
Delete	Administrator can (not their own account)	Never
Change password	Self-change or administrator reset	Self-change only

The Admin account is also exempt from the email-address requirement — it is the one sign-in name that is not an email, and it never receives the one-time email update prompt described in Chapter 2.

Note: Because `Admin` can never be locked out by account management, its password is the single most important credential on the DM. Set it to something strong, store it securely, and reserve the account for recovery and initial setup — give each administrator a personal Admin-level account for day-to-day work, so the Logons log shows who actually did what.

The Logons audit log

Every sign-in and every account change on the DM is recorded in the **Logons** log, the console's security audit trail. Each entry is time-stamped and includes the IP address the request came from. Recorded events include:

- **Successful sign-ins** — who signed in, from where.
- **Failed sign-ins, with the real reason.** The sign-in page deliberately shows only the generic "Invalid credentials", but the log records what actually happened — wrong password, unknown user name, or a disabled account — so an administrator can tell a mistyped password from someone probing for account names.
- **Password changes and resets** — successes and failures.
- **Account changes** — accounts created, renamed, and deleted, and profile updates, each attributed to the user who made the change.

Reviewing this log periodically is good practice: a run of failed sign-ins from an unfamiliar IP address is worth investigating, and the account-change entries answer "who created this account?" long after anyone remembers. The log is viewed with the system log viewer under **Logs** — see the Logs section of Chapter 12 for opening, searching, and exporting logs.

To recap the View-user experience of this chapter: a View user can open the Users page and see exactly one row — their own. From it they reach their own account page, where they can change their first and last name, their sign-in email, and their password. Everything else here — adding users, the Enable toggles, access levels, password resets, and deletion — is in the administrator's hands.

Chapter 12 — System Settings

The **Settings** area is where administrators look after the gateway itself rather than the equipment it monitors: the system clock, configuration backups, software updates, system logs, password policies, and the network interfaces the console answers on. Most of what this chapter covers is Administrator territory — View users can browse nearly every page described here, but the buttons that change things are reserved for Admin accounts. Each section below ends with a note on what View users see instead.

The Settings hub

Click **Settings** in the navigation tree to open the Settings hub — a dashboard of every system-level configuration area on the gateway.

- In **Cards** view, each child area (for example **Access Control**, **Backups**, **Communication Ports**, **Mappings**, **Network Interfaces**, **Policies**, **Time**, **Updates**) appears as a card showing its icon, name, and a short description. Hovering highlights the card; clicking it opens that area.
- In **Table** view, the same areas are listed in two columns, **Name** and **Description**. Rows are clickable.
- The view you choose is remembered as your preference across the console.
- The page refreshes automatically at your configured refresh interval, and the status dot in the top-right corner shows the result of the last refresh (blue while loading, green on success, red on failure).

System information (license, serial, and version)

Administrators see a **Properties** button at the top right of the hub. Click it to open the **Settings Properties** dialog, which is where the gateway's identity lives.

The dialog shows a mix of read-only system information and editable settings:

- **Service Status** — whether the gateway service is running.
- **System Serial Number** — the serial number of this gateway.
- **System Version** — the installed software version.
- **System Installation Date** — when the system was installed.
- **Security Key Status** — the license key installed on this gateway.
- **Console IP Addresses** — the list of IP addresses (with subnet masks) the console currently answers on. Useful when you need to reach the gateway from another address on the network.

Read-only fields are shown grayed. Click **Save** to write any changes, or **Close** to leave without saving. If a save is rejected, the reason appears inside the dialog; a value containing the `|` character is blocked before saving with the message "Property "<name>" cannot contain the | character."

Rebooting the gateway

The Settings Properties dialog also contains the **Reboot Gateway** switch. Rebooting is a two-step, deliberately confirmed action:

1. Open **Settings** → **Properties**.
2. Turn on the **Reboot Gateway** switch and click **Save**.
3. Instead of saving immediately, the **Reboot Gateway** confirmation dialog opens: "Reboot the gateway now?" with the explanation "This restarts the system. You will lose this connection — the page will wait while it restarts and return you to the sign-in screen automatically once the gateway is back online."
4. Click **Reboot Now** (red) to proceed, or **Cancel** to back out.

After you confirm, a full-page overlay with a spinner covers the console while the gateway restarts. The overlay reports its progress — "Waiting for the gateway to come back online... <n>s (this can take a few minutes on some hardware)" — and the console checks for the gateway automatically every few seconds, for up to 5 minutes.

When the gateway answers again, your browser is returned to the sign-in page automatically (the reboot ends your session, so you sign in again). If the gateway has not responded after 5 minutes, the overlay reports: "The gateway has not come back online yet. It may still be starting up — please refresh this page in a minute, or check the gateway."

Caution: Rebooting restarts the entire gateway, including monitoring, alarm processing, and protocol services. Anyone else signed in to the console loses their connection too.

Admin only: The **Properties** button — and with it the reboot flow — is only available to Administrators. View users see the Settings hub and can browse into its child areas, but have no Properties button.

Date & Time

Click **Settings** → **Time** to view the gateway clock settings. The page is a read-only summary; all changes go through its Properties dialog.

The table has three columns — **Property**, **Value**, and **Description** — and shows:

- **Enable** — whether time management is enabled (**Yes/No**).
- **Timezone** — the gateway's timezone setting.
- **DST** — whether Daylight Savings Time adjustment is enabled.
- **NTP** — whether the clock synchronizes with a network time server.
- **NTP Server** — the time server used (for example `time.windows.com`).
- **Date and Time** — the current date and time, continuously updated by the time service.

The footer note reads "Use Properties to edit time settings." for Administrators, and "Time settings shown in read-only mode." for View users. This page does not auto-refresh; revisit it (or reopen it from the tree) to see updated values.

Changing time settings

Administrators see a **Properties** button at the top right. Clicking it re-reads the current values first, so the dialog never shows a stale clock.

The **Modify Time Properties** dialog contains:

- **Name** — permanently disabled and always "Time". This is a fixed system area; it cannot be renamed or deleted.
- **Enable Time Management** — the master switch for the whole dialog.
- **Timezone Setting** — a drop-down of timezones.
- **Enable Daylight Savings Time.**
- **Synchronize with Network Time Server** (NTP) and **Network Time Server**.
- **Date** and **Time** — native date and time pickers for setting the clock manually. The date picker accepts dates from 2000-01-01 through 2099-12-31, and the time picker takes hours and minutes (seconds are always saved as zero). The pickers themselves refuse impossible values such as February 30th.

The fields gate each other in two tiers:

1. **Enable gates everything.** While **Enable Time Management** is off, every other field in the dialog is disabled. Turning it on enables them instantly; turning it off grays them again.
2. **NTP gates the manual clock.** While **Synchronize with Network Time Server** is on, the **Date** and **Time** pickers are disabled — the time server owns the clock, so a manually entered time would be overwritten. Timezone, DST, and the server name stay editable with NTP on.

In short: to set the clock by hand, turn Enable on and NTP off; to let a time server manage the clock, turn both on and fill in the server name.

The **Save** button starts disabled and only lights up once you have actually changed something; undoing all your edits disables it again. If a required picker is left empty, saving is blocked with "Date is required." or "Time is required." On a successful save the dialog closes and the page reloads with the new values.

Note: What View users see instead — the same read-only table, with no Properties button and the footer "Time settings shown in read-only mode."

Backups

The gateway can save its entire configuration to a backup file. Click **Settings** → **Backups** to see every backup the system holds.

Each backup file ends in a day-of-week number — `_0` for Sunday through `_6` for Saturday — so there is one slot per weekday, and a new backup taken on the same weekday overwrites that day's earlier file (its timestamp simply updates). A backup you start yourself with **Backup Now** is named `DT_Backup_<n>.bk` (for example, `DT_Backup_1.bk` for a Monday). You will also see automatic backups the system makes of its own components, named for the component — for example `CORE_SP01_Backup_<n>.bk` and `DM_SP01_Backup_<n>.bk`.

Administrators see two buttons on the left of the section header:

- **Backup Now** — take a backup immediately.
- **Delete Selected (n)** — red; disabled until at least one backup is checked. The count updates live as you select.

On the right are the **Properties** button (the backup schedule — see below) and the **Cards / Table** toggle. In Cards view each backup card carries a selection checkbox in its top-right corner; in Table view a leading checkbox column appears, with a select-all checkbox in the header. Your selection is kept across the page's auto-refresh and is shared between the two views. Clicking a backup itself (not its checkbox) opens that backup's detail page.

Backing up now

1. Click **Backup Now**. The **Create Backup** dialog asks "Create a backup now?" with the caution "This will save the current configuration."
2. Click **Backup Now** to confirm (or **Cancel**).
3. The dialog switches to an in-progress state — a spinner with "Backing up..." — and cannot be dismissed while the backup runs.
4. The console watches the backup list for up to 60 seconds. As soon as the new file appears (or the current day's existing file updates), the dialog closes and the list refreshes.

If the backup takes longer than 60 seconds, the dialog reports "Backup is taking longer than expected." / "Please refresh manually and check the backups list." with a **Close** button. This does not necessarily mean the backup failed — on slower hardware it may still complete. Close the dialog, wait a moment, and check the list.

The backup schedule

Backups are also written automatically on a schedule. Click the **Properties** button (gear) on the Backups page to edit it. The schedule dialog includes a **Time of Backup** field — a single time picker (HH:MM) for when the daily backup runs. Clearing the picker saves as 0:00 (midnight). Click **Save** to apply, **Close** to discard.

Viewing and downloading a backup

Click a backup in the list to open its detail page. The page shows a read-only **Property | Value | Description** table with details recorded by the backup service, such as the file size and when the backup

started and finished.

Administrators get two toolbar buttons:

- **Download** — saves the `.bk` file through your browser's normal download flow, using the backup's filename. Keep downloaded backups somewhere safe off the gateway; they are your recovery path if the gateway's configuration is ever lost.
- **Properties** — opens a read-only view of the backup's properties. Backup data is never editable, so there is no Save; the filename field is fixed too (renaming a backup would break its link to the file on disk). The dialog's footer holds a red **Delete** button and **Cancel**.

Deleting backups

You can delete a single backup from its detail page, or several at once from the list.

Single backup: open the backup, click **Properties**, then **Delete**. The confirmation asks "Delete <filename>?" with the caution "This action cannot be undone. The file is removed immediately." Confirm with **Delete** (red); the console then returns you to the Backups list.

Multiple backups:

1. On the Backups list, check each backup you want to remove — **Delete Selected (n)** enables and shows the count.
2. Click **Delete Selected (n)**. The **Delete Backups** dialog asks "Delete <n> backup(s)?" with the caution "This action cannot be undone. The files are removed immediately."
3. Click **Delete** (red) to confirm.

Caution: Deleting backups cannot be undone. If any file in the batch could not be deleted, it stays selected and an alert lists the affected filenames so you can retry.

Note: What View users see instead — the backups list and each backup's detail table, read-only: no Backup Now, no checkboxes, no Delete, no Download, no Properties.

Software Updates

Software updates arrive as `.upd` files supplied by Pythia Technologies. Click **Settings** → **Updates** to upload and apply them. The page subtitle summarizes the current state — for example "2 pending • 5 applied • 1 failed" — or "No updates uploaded yet." on a fresh system.

Uploading update files

Administrators see a large dashed upload zone at the top of the page: "Drag .upd files here", with a **Choose Files** button as an alternative to dragging (clicking anywhere in the zone also opens the file picker, and you can select several files at once).

Files upload one at a time, and a progress strip below the zone tracks each file: queued, then a spinner with "uploading...", then a green "uploaded" or a red error message. Files that are not `.upd` are skipped with "skipped (not .upd)". The strip clears itself a few seconds after the batch finishes, and the list refreshes to show the new files.

A freshly uploaded update appears in the list with status **PENDING**. Re-uploading a file with the same name replaces it and resets it to a fresh Pending state (any previous applied time or error is cleared).

Reading the list

Each update shows a status badge:

- **PENDING** (blue) — uploaded but not yet applied.
- **APPLIED** (gray) — applied successfully.
- **FAILED** (red) — the last apply attempt did not succeed.

The list is ordered so unfinished work comes first: Pending updates at the top, then Failed, then Applied, newest upload first within each group. Cards and table rows show the filename, status badge, upload time, and applied time (when present). Clicking an update opens its detail page; Administrators also get a selection checkbox on each entry.

Applying updates

1. Check one or more updates, then click **Apply Selected (n)**.
2. The **Apply Updates** dialog asks "Apply & update(s)?" with this caution: "Each update DLL runs with system privileges and may modify the registry, the database, or system files. This action cannot be undone -- there is no rollback."
3. Click **Apply** to confirm (or **Cancel**).
4. The dialog switches to a spinner — "Applying & update(s)..." — and cannot be dismissed while the updates run. Updates are applied one after another.
5. The console watches each selected update for up to 60 seconds, until every one has moved off Pending (to Applied or Failed). The dialog then closes and the list refreshes with the new badges.

If the batch takes longer than 60 seconds, the dialog reports "Updates are taking longer than expected." / "Check the list for current status -- some may have completed." Close it and watch the list; the badges update as the service finishes each file.

Caution: Applied updates cannot be rolled back. Take a backup (see the previous section) before applying updates, and apply them during a maintenance window when a service interruption is acceptable.

Update details, re-applying, and deleting

Click an update to open its detail page. The title carries the filename and its status badge, and the table shows **Status**, **Upload Time**, **Applied Time**, **Last Error** (only after a failed apply), and **Description**

(written by the update itself, once applied).

The toolbar follows the status:

- A **Pending** update offers **Apply Now** — same confirmation and caution as Apply Selected, for just this file.
- A **Failed** update offers **Re-Apply** — the confirmation asks "Re-apply <filename>?" and notes that the previous Last Error will be cleared on success or replaced on failure. Check **Last Error** first; it tells you why the previous attempt failed.
- An **Applied** update has no apply button.

To remove an update, open its **Properties** (a read-only view, like backups) and click **Delete**. The confirmation carries two cautions: "The .upd file will be removed from disk and the entry will disappear from the list." and "This does NOT roll back the effects of any previously-applied update." Confirm with **Delete** (red); the console returns you to the Updates list.

Note: What View users see instead — the update list and detail pages with badges and times, but no upload zone, no checkboxes, and no Apply, Re-Apply, or Delete buttons.

System Logs

The gateway keeps a set of system logs — one per service area (alarms, automations, console logins, updates, and so on). Click **Logs** in the navigation tree to see them all.

Each card shows a small round dot — green when logging is enabled for that log, gray when it is disabled — plus the log's name and description. In Table view the columns are **Log Name**, **Enabled** (a green "✓ Yes" or a muted "— No"), and **Description**. Click a log to open it in the log viewer. This dashboard is read-only for everyone; enabling and disabling happens inside the viewer.

The log viewer

The viewer shows the newest 500 entries of the selected log, newest first, in two sortable columns: **Time** and **Text**. Click a column header to sort by it; click again to flip the direction. Above the table:

- **Search** — filters the entries as you type, matching against both time and text. While a filter is active, a clear (X) button appears and the footer shows "<matches> of <total> entries". If nothing matches: "No entries match <filter>".
- **Logging Enabled / Logging Disabled** — a toggle that turns logging for this log on or off. It flips immediately; if the gateway rejects the change (for example, you are not an Administrator), the switch flips back and an alert explains why.
- **Refresh** — reloads the entries. The viewer does not refresh on its own.
- **Export: PDF / Excel** — see below.
- **Clear Log** — permanently empties the log.

The gateway also prunes very old entries automatically in the background, so logs do not grow without bound.

Clearing a log

Click **Clear Log** and the browser asks: "Delete all entries from the <log name> log? This cannot be undone." Click **OK** to delete every entry and reload the (now empty) viewer.

Caution: Clearing a log removes all of its entries permanently. Export the log first if you may need the history.

Exporting a log

PDF produces a landscape report titled "<log name> Log" with a Pythia Technologies header, the export timestamp, a note of any active search filter, and the entries in a striped Time/Text table. **Excel** produces a workbook with the same header information and columns. Both exports contain **only the currently filtered view, in the current sort order** — so you can search for what matters first, then export exactly that. Files are named <log name>_Log_YYYYMMDD_HHMM.pdf / .xlsx.

Admin only: The **Logging Enabled** toggle and **Clear Log** require Administrator access. The viewer shows these controls to everyone, but a View user who tries them gets a failure alert containing "0003 - Access denied" and nothing changes. Search, Refresh, and both exports work for all users.

Password Policies

Password policies let you enforce organization-wide rules on console passwords. Click **Settings** → **Policies** to open the policy categories, then click **Passwords**.

The **Passwords** category contains five policies:

Policy	What it enforces
Admin Maximum Age	Maximum age of Admin password in days
History	Block reuse of the last number of previous passwords
Lockout	Lock account after number of consecutive failed sign-ins
Minimum Length	Minimum password length in characters
User Maximum Age	Maximum age of user password in days

Note: All password policies ship **disabled**. Nothing is enforced until an Administrator turns a policy on.

Each row has an inline **Enable** toggle, so you can switch a policy on or off right from the list — the switch flips immediately and the policy's shield icon recolors (green = enabled, gray = disabled). If the save fails, the switch reverts and an error banner appears above the table. The footer reminds you of the split: "Toggle Enable inline. Click a policy name to edit its value."

Editing a policy value

Click a policy's name to open its editor. The page shows:

- A read-only card with the policy's **Policy** name and **Description**.
- An editable card with the **Enable** toggle and the **Value** — a whole-number input with the hint "Allowed range: <Minimum> – <Maximum>" — plus the read-only **Minimum** and **Maximum** bounds themselves.

The value is checked as you type: an empty or non-numeric entry shows "Value must be an integer.", and an out-of-range entry shows "Value must be between <min> and <max>." — either one turns the input red and disables **Save**. Click **Save** to apply (the page confirms with a green "Saved." and the header updates), or **Cancel** to put the form back the way it was. If you save without changing anything, the page simply says "No changes to save."

The policy names, descriptions, and Minimum/Maximum bounds are fixed at install time; only Enable and Value are yours to change.

Note: What View users see instead — the same category list and editor pages with every toggle and input disabled, and the footers "Read-only view. Sign in as an admin to enable / disable policies." and "Read-only view. Sign in as an admin to edit."

Network Interfaces

Click **Settings** → **Network Interfaces**, then an interface number (**0**, **1**, ...) to configure how the gateway connects to your network. Each interface gets its own page.

The adapter card

The top card is read-only, live information discovered from Windows:

- **Name** — the Windows adapter name.
- **MAC** — the adapter's hardware address.
- **Status** — shown green and bold when **Connected**, dimmed otherwise.
- **Link Speed** — the negotiated speed (for example 1 Gbps).
- **Current IP** — the address the adapter has right now.

Values the gateway could not read show as a dash (—).

The configuration card

Below it, Administrators can edit the interface's configuration:

- **Address Mode** — a toggle labeled **Obtain an IP address automatically (DHCP)**. This toggle appears **only on interface 0**; every other interface is always static. In DHCP mode the IP table is replaced by the note "IP address, subnet mask and gateway are assigned automatically by DHCP."
- **IP Addresses** (static mode) — a table of the interface's addresses with columns **IP Address**, **Subnet Mask**, and **Gateway (optional)**. Each row has an × button to remove it, and the dashed + **Add IP** button below adds another row (the gateway stores up to 16 addresses per interface). Rows left completely empty are ignored when you apply.
- **DNS Server** — the DNS server address; in DHCP mode the label adds "(blank = automatic)".
- **Cancel** and **Apply & Restart** — both stay disabled until you actually change something. **Cancel** throws away your edits, reloads the saved configuration, and confirms with a green "Changes discarded."

Applying changes (Apply & Restart)

Network changes take effect only when you apply them — and applying **reboots the gateway**.

1. Make your edits, then click **Apply & Restart** (red).
2. Your entries are checked first. Problems appear in red below the buttons — for example "Row 1: invalid subnet mask.", "DNS must be a valid IPv4 address.", or "Add at least one IP address (or switch to DHCP)." Fix them and click again.
3. The **Apply & Restart** confirmation opens: "Apply the network configuration and restart the system?" with the caution "This pushes all interface changes to Windows and reboots the system. If you changed the IP of the interface you're connected through, you will lose this connection and must reconnect at the new address."
4. Click **Apply & Restart** to confirm (or **Cancel**).

Caution: Apply & Restart applies the configuration of **all** interfaces and reboots the gateway. Monitoring stops during the restart, and every signed-in user loses their connection. Double-check the addresses before confirming — a wrong static IP can leave the gateway unreachable from your workstation.

What happens next depends on your configuration:

- **Static address set:** the dialog shows a spinner with "The system is restarting — this usually takes about a minute. You'll be taken to the new address automatically once it's back online." and a clickable link, "New address: <http://192.0.2.10/>". The console then watches that address; once it sees the gateway go down and come back up, it confirms "✓ The system is back online. Continuing to <http://192.0.2.10/>..." and takes your browser to the new address automatically a moment later. (Clicking **Close** stops the automatic redirect if you would rather navigate yourself.)

- **DHCP / no static address:** the dialog shows "The system is restarting — this usually takes about a minute. Reconnect once it's back online." — the console cannot predict a DHCP-assigned address, so reconnect manually once the gateway is up.

Every network save and apply is recorded in the system logs with the user and address that made the change.

Note: What View users see instead — the full interface page with the adapter card and current configuration visible, but every input disabled and no Cancel or Apply & Restart buttons.

SMTP Relay

The gateway sends its own email — alarm notifications and heartbeat messages — through a small built-in mail relay that runs as a system service — enabled by default, and replacing the older IIS SMTP relay used on earlier gateways. The **SMTP Relay** page is where you watch that service work and configure how it delivers mail. Every message the gateway generates is handed to the relay, which either delivers it straight to the recipient's mail servers or forwards it through a **smart host** (your organization's mail server) on the way out. Because all of the gateway's outbound email flows through this one service, the dashboard here is the fastest place to confirm that notifications are actually leaving the building — and to see why they are not, if delivery has stalled.

The page opens on a live dashboard. The header shows the service status as a colored dot and a one-line summary beneath the title — for example the running state, the delivery mode, and the **HELO** name the relay announces to other mail servers. **Updated: HH:MM:SS** in the top-right corner records the last successful refresh, and the dot beside it turns blue while loading, green when the service is listening normally, and red if the service cannot be reached or is not running. The whole dashboard refreshes on its own every 5 seconds, so you can leave it open and watch mail move.

If the relay service is not present on this gateway, the page reports "SMTP Relay service not installed" and there is nothing to configure.

Reading the dashboard

A row of KPI cards across the top summarizes the service at a glance:

- **Queue** — how many messages are waiting to be delivered right now, with a small bar that fills as the queue grows and turns red once the queue reaches the alarm depth you set. Below it, the queue's size in KB and the configured message limit.
- **Delivered** — the running total of messages delivered since the counters were last reset, with the time of the most recent delivery.
- **Failures** — delivery problems, broken out beneath as "<n> failed · <n> expired" (messages the relay could not deliver, plus messages that aged out of the queue before they could be delivered).
- **Connections** — the number of inbound sessions the relay has accepted.

- **Accepted** — messages accepted and queued for delivery.
- **Rejected (Access)** — connections turned away because the peer is not on the allow list (see below); the sub-line shows the last address that was denied.
- **Rejected (Full)** — messages refused because the queue was at capacity.
- **Rejected (Size)** — messages refused for exceeding the message size limit.

Under the cards, a **Counters since <date>** line tells you how far back these totals reach. If the service has recorded a delivery problem, a red **Last Error** strip appears with the reason; when everything is healthy, the strip is hidden.

Below the KPIs are two panels:

- **Message Queue** — the messages currently waiting, one row each, with columns **Message**, **From**, **To**, **Age**, **Tries**, **Next Try**, and **Last Error**. Delivery is normally immediate, so this panel is usually empty; when it is, it reads "Queue is empty — delivery is normally immediate; messages appear here only while delivery is pending or retrying." A message only lingers here while the relay is retrying a recipient that is temporarily unreachable.
- **Recent Activity** — a rolling feed of the last several events, newest first, each tagged **QUE** (queued), **OK** (delivered), **FAIL** (failed), or **EXP** (expired). Before any mail has moved it reads "No activity yet."

Everything described so far is visible to every signed-in user. The buttons that change the service — **Send Test**, **Reset**, and **Properties** — appear only for Administrators.

Sending a test message

The quickest way to prove the whole path works — from the relay, through your smart host or the recipient's mail servers, to a real inbox — is to send a test message.

1. In the box marked **Send test email to...** at the top left, type a destination address, for example `user@example.com`.
2. Click **Send Test**.
3. The button reports "Sending...", then "Queued." once the relay accepts the message. Watch the **Recent Activity** feed for it to move from **QUE** to **OK**, and check the destination inbox to confirm it actually arrived.

If you leave the address empty or type something without an @, the relay answers "A valid destination address is required" and nothing is sent. If the relay accepts the message but the onward delivery is refused — by your smart host or the recipient's server — you will see "Relay refused the message: <reason>", where the reason is the text the receiving server gave back. A test that leaves the queue and lands in the inbox confirms the entire delivery chain, not just the relay itself.

Note: The test message is sent from the same address your real alarm mail uses, so it exercises the same acceptance and anti-spam policies at your smart host or the destination — a passing test is a genuine dress rehearsal for a real alarm.

Resetting the counters

Click **Reset** to zero all of the counters — Delivered, Failures, Connections, and the rest — and restamp the **Counters since** date to now. The button confirms with "Reset." once done. This is handy after you have finished troubleshooting and want a clean baseline. It clears the running totals and the recent-activity history; it does not touch any queued mail or your configuration.

Caution: Resetting the counters cannot be undone. The delivery, failure, and rejection totals collected since the last reset are cleared permanently. If you need those figures for a record, note them before resetting.

Configuring the relay

All of the relay's settings live behind the **Properties** button, which opens the **SMTP Relay Properties** dialog. This keeps the dashboard uncluttered while giving you full control of how mail is delivered. Changes take effect within a few seconds of saving — the service does not need to be restarted.

The dialog's configuration grid holds:

- **Smart Host** and **Smart Host Port** — the mail server to forward all outbound mail through, and its port. Leave the smart host blank to deliver directly to each recipient's own mail servers instead (the field shows **Direct MX (no smart host)** as a reminder when empty).
- **From Address** — the sender address on the gateway's mail. Use an address your mail system will accept as a legitimate sender.
- **HELO Name** — the name the relay announces when it connects to another mail server.
- **Message Limit MB** — the largest single message the relay will accept; anything larger is counted under **Rejected (Size)**.
- **Queue Alarm Depth** — the queue size at which the **Queue** card turns red. Set it to 0 to disable the queue alarm.
- **Queue Limit** and **Queue Limit MB** — the most messages, and the most total megabytes, the queue may hold before new mail is refused (counted under **Rejected (Full)**).
- **Expire Days** — how long the relay keeps retrying an undeliverable message before giving up on it.
- **Keep Undeliverable (Badmail)** — whether to keep messages that expire undelivered (**On — keep in Badmail**) or discard them (**Off — delete on expiry**), and **Badmail Keep Days** for how long kept copies are retained.
- **Bind Address(es)** — which of this server's own network addresses the relay listens on, shown as a checklist of the gateway's addresses. Leave every box unchecked to listen on all interfaces.

Each numeric field is validated when you save. A value that is not a positive whole number (or, for the port, not between 1 and 65535) is rejected with "<field> value is invalid", and nothing is written until every field is valid — so a single bad entry can never leave a half-applied configuration behind. Clearing a field returns that setting to its built-in default.

Click **Save Settings** to apply; the dialog reports "Saving...", then "Saved." Click **Close** to leave the dialog.

The connection Allow List

The relay listens on the network so that the gateway's own services can hand it mail, but it does not accept connections from just anywhere. By default only the gateway itself (the loopback address, `127.0.0.1`) may connect. If you need another host to submit mail through this relay, add it to the **Connection Allow List** at the bottom of the Properties dialog.

- To add a peer, type its address in the **IP, range, or hostname** box, optionally add a note in **Description (optional)**, and click **Add**. You can enter a single IPv4 address (for example `192.0.2.40`), a CIDR range (for example `192.0.2.0/24`), or a hostname. The list confirms with "Added." An entry that is not one of those forms is refused with "Not a valid IPv4 address, CIDR range, or hostname", and an empty entry with "An entry (IP, range, or hostname) is required". Adding an address that is already listed simply updates its description.
- To remove a peer, click **Remove** on its row; the list confirms with "Removed." When the list is empty it reads "No entries — only 127.0.0.1 may connect."

The loopback address is always allowed and is not shown in the list — you cannot remove the gateway's own access to its relay. Any host that is neither loopback nor on the allow list is refused at connection time with an SMTP **554** rejection, and that attempt is counted under **Rejected (Access)** on the dashboard with the peer's address shown as the last denied connection.

Caution: Removing a peer from the allow list takes effect within a few seconds, and that host's connections are refused with a 554 rejection from then on. If that host relies on this relay to send mail, its messages will stop going out — confirm the address is genuinely unwanted before removing it.

Admin only: **Send Test**, **Reset**, **Properties**, and everything inside the Properties dialog — the relay configuration and the Connection Allow List — require Administrator access. The gateway rejects these actions from other accounts even if the controls were reached some other way.

Note: What View users see instead — the full live dashboard (status, all KPI cards, the message queue, and the recent-activity feed, refreshing every 5 seconds) exactly as an Administrator sees it, but with no **Send Test** box, no **Reset** button, and no **Properties** button. View users can watch the relay but cannot send a test, reset the counters, or change any setting.

Chapter 13 — Integrations: SNMP, Modbus and BACnet

The DM does more than monitor equipment — it re-publishes everything it monitors, so a building-management system, a network management platform or any other supervisory system can read the same points over the protocol it speaks natively. This chapter covers the three protocol mappings (SNMP, Modbus/TCP and BACnet), the point lists you can export from each, the ready-made ALC / WebCTRL integration URLs, and the **Views** page that shows the health of the background services doing the collecting.

Note: This chapter is written for **integrators** — the person pointing an external BMS or NMS at the DM. You will want it open alongside the configuration tool of the system you are connecting. Day-to-day operators can safely skip ahead; nothing here changes how the console itself monitors equipment.

How the mappings work

Every alarm, data point and control you configured under **Locations** is automatically assigned a place in three parallel address spaces:

- an **SNMP OID** under the Pythia Technologies enterprise prefix,
- a **Modbus/TCP register** (or coil) on a per-unit device id, and
- a **BACnet object instance** on a per-unit BACnet device.

You never build these mappings by hand. The DM's mapping service generates and maintains them, deriving all three addresses for a point from one internal mapping ID — which is why a point keeps the same address even if you rename its unit, and why the reset action described later rennumbers all three protocols at once.

The mappings live under **Settings** → **Mappings** in the navigation tree. The **Device Mappings** page lists the three protocols as clickable cards (or table rows, if you prefer Table view — see Chapter 3):

Admin only: The **Properties** button on this page is the home of the global **Reset Mapping IDs** action, covered at the end of the mappings sections. View users see the three protocol cards without a toolbar.

Click a protocol to open its mapping root. The three roots share one layout: a header with the mapping name and a unit count, a live search bar, an export bar, and a sortable table with one row per mapped unit. Click any column header to sort (the active column is highlighted with a ▲/▼ arrow), and click a row to drill down into that unit's points.

Note: Unlike most console pages, the three mapping roots and their drill-downs do **not** auto-refresh — the mapping changes only when your configuration does. Use the **Refresh** button in the header to re-read the page.

The Modbus/TCP mapping

Open **Settings** → **Mappings** → **Modbus/TCP**. Each mapped unit appears as one Modbus device:

Column	Meaning
Device Id	The unit's Modbus device (slave) id. Stored as a three-digit number; displayed without leading zeros, so device 007 shows as 7.
Mapped Unit	The source unit, as its full path — for example \Locations\Sunbury Office\Power Units\UPS-21A.
Points	How many registers and coils are mapped under that device.

The search bar filters by device id and mapped unit as you type; a × button clears it. If nothing matches you see "No units match" with your search term, and an unconfigured system shows "No units mapped."

Click a row to drill into the device and walk down to its individual points. Points fall into three register ranges: **coils** (addresses below 10000) carry the unit's alarms, the **30000 range** carries its data readings, and the **40000 range** carries its controls. Writable setpoint registers carry an **Enable** property, edited on the point's own properties page — or all at once with **Enable All R/W**, described below.

Integer vs Float: the mapping Format

The one thing you can configure on the Modbus mapping is how analog data values are encoded. Click **Properties** (admin only) on the Modbus/TCP root and set the **Format** dropdown:

- **Float HiLo** (the default) — each analog value occupies a 32-bit IEEE float across a **pair** of consecutive registers, high word first.
- **Float LoHi** — the same float pair with the word order reversed, for masters that expect low word first.
- **Integer** — each analog value is a single 16-bit scaled integer register.

The dialog carries a standing advisory, and it means what it says:

Note: "Changing the format rebuilds the Modbus register map. After saving, allow 1-2 minutes for the service to update the registers before exporting." Exporting — or polling — during the rebuild window can catch the map mid-change.

Word order matters only for the Float formats: if your master reads garbage where a sensible reading should be, the two words of the pair are almost certainly swapped — switch between **Float HiLo** and **Float LoHi**, or flip the equivalent word-order option on the master. (For ALC / WebCTRL specifically, see the FLOAT note in the URL section later in this chapter.)

The SNMP mapping

Open **Settings** → **Mappings** → **SNMP**. Every point the DM publishes lives under the Pythia Technologies enterprise OID prefix **1.3.6.1.4.1.38801** — the mapping root sits one level below it (for example **1.3.6.1.4.1.38801.4**), and each mapped unit occupies the next arc down. Because units differ only in that final arc, the root table sorts by OID numerically:

Column	Meaning
OID	The unit's complete dotted OID — the base OID plus the unit id.
Unit Id	The unit's index within the SNMP table.
Mapped Unit	The source unit path.

The search bar here filters by id, unit and OID ("Search units / OIDs...").

Drilling into OIDs

Click a unit row to open the SNMP node page for that unit, and keep clicking to descend — unit, then its Alarms / Data / Controls branches, then the points. The design goal of these pages is that you never lose sight of the full address: the header shows an **OID** chip with the complete dotted OID of the level you are on, and the table's first column is the **Complete OID** of each child, ready to paste into your NMS.

Alongside the Complete OID, each row shows:

- **Enable** — **Yes** / **No**, shown only on points that can be written (setpoints); blank everywhere else.
- **Type** — **Readonly** or **Read/Write** on leaf points; blank on branches.
- **Description** — the point's name.

Clicking a branch row descends another level; clicking a leaf point opens its properties page, where an administrator can switch an individual setpoint's **Enable**. The drill-down pages themselves are identical for Admin and View users — no editing happens there.

The SNMP mapping has its own **Format** property (in **Properties** on the SNMP root, admin only): **String** (the default — values are reported exactly as displayed) or **Integer** (a legacy scaled-integer encoding for older managers).

The BACnet mapping

Open **Settings** → **Mappings** → **BACnet**. Like Modbus, the layout is flat — one BACnet device per mapped unit:

Column	Meaning
Device Id	The unit's BACnet device id (three digits, shown without leading zeros).
Mapped Unit	The source unit path.
Objects	How many BACnet objects the device exposes.

Every object the DM publishes is a BACnet **Analog Value** — there is no encoding format to choose for BACnet. The root's **Properties** dialog (admin only) still exists for the container's other settings, including the **BACnet Instance** device id the DM server itself answers as.

Object instances — the device-id math

Click a device row to list its objects. Internally, objects are numbered from a low base per point type (1001, 3001, 4001, ...), and the **complete instance** your BMS must address is:

complete instance = device id × 10000 + low instance

So object 3001 on device 1 is instance **13001**; the same object on device 12 would be **123001**. You never have to do this arithmetic yourself — the **Instance** column on the device page shows the complete, ready-to-use number:

The remaining columns match the SNMP drill-down: **Enable** (Yes/No on writable setpoints only), **Type** (Readonly / Read/Write) and **Description**. Clicking an object opens its properties page.

Enable All R/W

By default, writable setpoints in the mappings are individually enabled — a deliberate safety posture, since an enabled read/write point lets an *external* system command your equipment. When you are ready to hand write access to the integrated system, you can enable every setpoint in a protocol at once instead of visiting each point.

Admin only: The green **Enable All R/W** button appears at the right of the export bar on each of the three mapping roots, for administrator accounts only.

1. On the mapping root, click **Enable All R/W**.
2. A browser confirmation asks: "Enable ALL read/write (setpoint) points for this protocol?" Click **OK** to proceed.
3. On success an alert reports the result — "Enabled 12 read/write point(s)." — and the page reloads. The action covers every setpoint under that protocol's mapping, however deep.

The action applies to one protocol at a time; run it on each root you are integrating. Individual points can still be switched back off afterward on their properties pages.

Exporting the point lists

Each mapping root has an export bar — **Export:** followed by a row of format buttons. Exports are generated in your browser and download immediately; nothing opens in a new window (except **ALC URL**, covered next). Every export contains the full flattened point list for that protocol, sorted numerically by address.

Format	Modbus/TCP	SNMP	BACnet	Notes
CSV	✓	✓	✓	Plain comma-separated point list
PDF	✓	✓	✓	Formatted, printable
Excel	✓	✓	✓	.xlsx workbook
HTML	✓	✓	✓	Self-contained web page
JSON	✓	✓	✓	For scripted consumers
MIB	—	✓	—	A MIB file your NMS can compile, so points browse by name
EDE	—	—	✓	ASHRAE Engineering Data Exchange point list (.csv), the standard BACnet hand-off document

Files are named after the mapping and time-stamped — for example

Modbus_TCP_Mapping_20260705_0910.csv, SNMP_Mapping_20260705_0910.mib, or

BACnet_Mapping_20260705_0910_EDE.csv. The BACnet exports also carry the DM's vendor id (617) and its BACnet device instance, which EDE-aware tools read directly.

The first PDF or Excel export in a session pauses a moment while the generator library loads; if it cannot load you get an alert ("PDF library failed to load." / "Excel library failed to load.") — check the connection and try again.

Note: Export a fresh copy after any change that rennumbers or re-formats the mapping — adding or deleting units, changing the Modbus **Format**, or running **Reset Mapping IDs**. An export made before the change describes a map that no longer exists.

Reset Mapping IDs

If a site's mappings have accumulated gaps — units added and deleted over the years leave unused ids behind — an administrator can renumber everything from scratch. Because all three protocols derive from the same internal ids, this is a single global action, not a per-protocol one.

Admin only: Reset Mapping IDs lives inside the **Properties** dialog of the **Device Mappings** parent page (Settings → Mappings) — a red **Reset Mapping IDs** button in the bottom-left corner of the dialog.

1. On the **Device Mappings** page, click **Properties**.
2. Click the red **Reset Mapping IDs** button. The Properties dialog closes and a dedicated confirmation opens, asking "Reset all mapping IDs and start from scratch?" with a caution box spelling out the consequences.
1. Click the red **Reset Mapping IDs** button in the confirmation to run it. The button disables while the reset runs; on success the dialog reports "Cleared 214 mapping ID(s)." (with your actual count) and reminds you that "Renumbering is in progress — allow up to ~2 minutes for the SNMP, Modbus and BACnet mappings to rebuild." Click **Close** to finish.

Caution: As the confirmation itself warns: "This renumbers every SNMP, Modbus and BACnet mapping at once — OID, register and BACnet instance numbers will all change. Any connected clients must be re-pointed afterward. Renumbering completes within ~2 minutes. This cannot be undone." Every external system polling the DM — every NMS OID, every BMS register and instance reference — must be re-mapped against the new numbers. Plan this with the owners of those systems before you click, and export fresh point lists (and new ALC URLs) once the rebuild settles.

ALC / WebCTRL integration URLs

For ALC / WebCTRL systems, the DM goes one step further than a point list: it writes the integration URLs for you. Each of the three mapping roots has an **ALC URL** button at the end of its export bar; clicking it opens a dedicated page **in a new browser window** listing one ready-to-paste URL per point. The page is read-and-copy only — Admin and View users see the same thing.

Setting the parameters

Some pieces of each URL depend on your installation, so the toolbar offers selectors — change one and every URL on the page updates instantly:

- **SNMP** — an **SNMP OID Table Index** dropdown, **S1** through **S80**. WebCTRL abbreviates long OIDs through its SNMP OID Table: you register the DM's enterprise prefix under an index, then every point URL references that index. A note box above the table tells you exactly what to enter — "Enter this OID prefix into index S5 of the SNMP OID Table in WebCTRL:" followed by `S5 = 1.3.6.1.4.1.38801` — and updates as you change the dropdown.
- **BACnet** — a **Server IP Address** dropdown listing the DM's own console addresses; pick the one your BMS reaches the DM on.
- **Modbus/TCP** — the same **Server IP Address** dropdown only. The Modbus unit address in each URL is filled in automatically from the mapped unit's own device id, so there is no address selector

to set.

The URL grammar

Protocol	Pattern	Example
SNMP	snmp://1/s<n>/<OID with the 1.3.6.1.4.1.38801 prefix removed, keeping the leading dot>	OID 1.3.6.1.4.1.38801.4.1.1.1 at index S5 → snmp://1/s5/.4.1.1.1
BACnet	bacnet://0:<server-ip>/2:<complete instance> (0 = network, 2 = Analog Value)	bacnet://0:192.0.2.10/2:13001
Modbus/TCP	mtcpip://<type>/<register>/<modbus-address>/<server-ip>	mtcpip://FLOAT/30001/7/192.0.2.10 (here 7 is the mapped unit's Modbus device id)

For Modbus, the type token follows the mapping's **Format** property, and the page's note box states the current rules — coils (alarms) are always `Do`, controls are always `uint`, and 30000-range data points are `uint` under Integer format or `FLOAT` under a Float format. The list is also collapsed to one row per *logical* point: under a Float format only the first register of each 32-bit pair is listed (the URL addresses the whole pair), and in the 40000 range only the writable setpoint register is listed, not its read-back echo. The `<modbus-address>` segment of each URL is the mapped unit's own Modbus device id, so different units' URLs carry different addresses — there is no single address to choose for the page.

Note: When the mapping Format is a Float variant, the note box adds: "For FLOAT points, ensure the word order in ALC matches the driver's 'Reverse order of High and Low Words for FLOAT Data Type' setting on the Modbus TCP/IP property page." If FLOAT points read as nonsense in WebCTRL, this word-order mismatch is the first thing to check.

Copying and exporting

- Each row has a **Copy** button that puts that one URL on the clipboard (the button flashes "Copied").
- **Copy All URLs** copies the whole visible list at once, flashing "Copied 96" with the count.
- The **Search** box filters by URL, unit, type or description — and the filter applies to **Copy All** and the exports too, so you can copy or export just one unit's points. The subtitle switches to "12 of 96 points" while a filter is active.
- The **Export:** buttons produce **CSV**, **PDF** and **Excel** files named like `Modbus_TCP_ALC_URLs_20260705_0910.xlsx`. Every export embeds the protocol note lines, and, when a filter is active, a line recording it — `Filtered by search: "UPS-21A"` — so a partial list can never masquerade as the full one.

The Views page — collector health at a glance

The last stop for an integrator is **Views**, at the top level of the navigation tree. Views are derived-data hierarchies (fleet roll-ups such as branch-circuit monitoring — see Chapter 6) built by dedicated background collector programs, and this page doubles as the health dashboard for those collectors.

Each configured view gets a card:

- A **status pill** — **Normal** (green), **Error** (red) or **Disabled** (gray) — with a matching colored edge on the card.
- A **collector line** — a green glowing dot with "Collector running (PID 4812)" when the collector process is up, or a gray dot with "Collector stopped".
- A **meta grid** with the collector's details: **Collector** (the program name, e.g. `PTA_VW_DFBCM`), **Poll Rate** (e.g. `30s`), **Started**, **Stopped** and **Runtime**.

A summary strip above the cards totals the fleet — `Views: 3 Running: 2 Stopped: 1 Errors: 0` — so one glance tells you whether the derived data feeding your dashboards (and your exports) is current. The page auto-refreshes every 5 seconds, faster than the normal list rate, so a collector starting or stopping shows almost immediately. Clicking a card drills into that view's data hierarchy in the tree.

On a system with no views configured, the page simply says so:

If a card shows **Error**, or **Normal** with the collector stopped, the view's data has stopped updating — restart the DM services or contact support before trusting downstream numbers.

What View users see

View users can read everything in this chapter — the Device Mappings pages, every drill-down with its complete OIDs, registers and instances, the ALC URL windows including Copy and the exports, and the Views page. What they cannot do is change anything: the **Properties** and **Enable All R/W** buttons on the mapping roots, and the **Properties** / **Reset Mapping IDs** action on the Device Mappings parent, appear only for administrators. This makes a View account a good fit for the integrator on the other side: they can look up addresses, copy URLs and export point lists without being able to renumber the map or enable write access.

Part 2 — DM (Data Monitoring)

The DM is the core platform itself. Part 1 is its complete guide.

DM — Data Monitoring

The **DM** is the core platform itself — the full data-monitoring product. Everything in **Part 1 — Core Platform** describes the DM directly: building your **Locations** → **Groups** → **Units** tree, adding data points and controls, configuring per-unit alarms, setting up automations and heartbeats, trending values over time, managing users, adjusting system settings, and exposing your points to other systems over SNMP, Modbus, and BACnet.

There are no DM-only features beyond the Core Platform. If you are running a DM, **Part 1 is your complete guide** — the rest of this document (Parts 3 and 4) applies only if you also run an SPDM/SPIGM or a SPAC4.

Note: Some DM installations include a wired **Smart Panel** even though the product is a plain DM. If your navigation tree has a **Smart Panel** node, the Smart Panel material in **Part 3** — including how to bring panel inputs into units with the DIRS driver — applies to you as well.

Part 3 — SPDM and SPIGM (Smart Panel Data Monitor)

SPDM (and the functionally identical SPIGM) is a DM built around a wired Smart Panel. The Smart Panel page is covered first, then how to bring its inputs into units with the DIRS driver. Everything else is the Core Platform.

Chapter 10 — Smart Panel

The **Smart Panel** page is a live monitor for the Pythia CM06P panel — a hardware I/O module whose inputs the DM reads continuously over a direct serial (Modbus RTU) connection. If your site includes a CM06P, this page is where you watch its input values in real time, map new module inputs as data points, and configure the panel itself. There is exactly one Smart Panel per DM: the page always shows that panel, and the panel cannot be deleted from the console.

Opening the Smart Panel page

Click the **Smart Panel** node in the navigation tree. The tree highlights the node and, beneath it, lists the panel's configured data points as child entries.

The page header shows, from left to right:

- **The panel name** — "Smart Panel" by default; an administrator can rename it through **Properties**.
- **The panel status badge** — a green-bordered badge (for example **Normal**) when the DM is communicating with the panel normally, a red-bordered badge when the panel is in communications alarm. Before the first load, or if the state is unknown, the badge shows a muted "--". The badge text is the panel's current status message.
- **The point count** — a subtitle such as "8 data points".

At the top right, the usual status dot and "Updated: HH:MM:SS" stamp (chapter 3) confirm the page is polling. The page refreshes at the system auto-refresh rate; values, communication states, the status badge and the warning banner all update live without any action on your part.

The data point table

Each row in the table is one configured input on the panel:

Column	What it shows
Data	The point's name (for example Input 01 or Temperature 1). For administrators the name is a link that opens the point for editing.
Description	Free-text description of the input — what is physically wired to it (for example "Input 1" or "(F)" for a Fahrenheit temperature).
Register	The Modbus register on the module that this point reads, shown in monospace (for example <code>1001</code>).
Type	The data format the register is read as (for example INT for a whole-number input, FLOAT32HILO for a 32-bit floating-point value such as a temperature).
Value	The value from the last successful read, shown in monospace.
Enable	Yes when the point is being polled; a muted italic No when it is disabled.

The **Data** column is also color-coded so you can read the panel's health at a glance:

- **Normal text** — the point read successfully on the last poll.
- **Red** — the point failed its last read.
- **Muted gray** — the point is disabled. A disabled point always shows gray, even if its last read also failed.

If no inputs have been mapped yet, the table is replaced by the message "No data points configured. Use Add to map a module input."

If the page itself cannot load, a red **Error** panel appears with the reason (or "Could not load Smart Panel." if none is available), and the subtitle changes to "Error loading panel".

When communications fail

When one or more points fail to read on a poll, a red warning banner appears above the table:

>  "Communications issue — N data point(s) failed to read on the last poll"

The count tells you how many points are affected, and the failing points' names turn red in the **Data** column. If the whole panel stops responding, the status badge in the header turns red as well.

A single red point usually indicates a problem with that mapping — most often a register number that does not exist on the module. Check the point's **Modbus Register** value (see "Editing a data point" below). If every point is red and the badge is in alarm, the DM cannot reach the panel at all — check the serial connection and the panel's communication settings in **Properties**.

Note: The banner clears on its own as soon as the affected points read successfully again. There is nothing to acknowledge or reset on this page.

Configuring the panel

Admin only: The toolbar buttons in this section appear only for Administrator accounts.

Administrators see two solid blue buttons above the table: **Add** (plus icon — "Add data point") and **Properties** (gear icon — "Panel settings").

Click **Properties** to open the panel configuration dialog (**Smart Panel Properties**). This is the standard Properties dialog described in chapter 3, with these settings:

- **Enable/Silence Buzzer** — turns the panel's audible buzzer on, or silences it.
- **Alarm Delay in Seconds** — how long a condition must persist before the panel raises an alarm.
- **Alarm Category** — a checklist (**Category 01** and up) choosing which alarm categories apply to the panel; the checkbox at the top of the list selects or clears them all at once.
- **Module Network** — the serial (COM) port the module is connected to, for example **COM5**.
- **Module Modbus Address** — the module's Modbus address on that port.

Make your changes and click **Save**; the dialog closes and the page refreshes with the new settings. Unlike other containers' Properties dialogs, this one has no **Delete** button — the Smart Panel is a permanent, single item and cannot be removed.

Adding a data point

To map a new module input:

1. Click **Add**. The **Add Data Point** dialog opens.
2. Enter a name for the point, its description, the **Modbus Register** it reads, and its **Type**. (The live **Value** is read-only and does not appear on the form.)
3. Click **Save**. While the dialog is loading or saving, **Save** is grayed out.

When the point is created, the dialog closes and the table refreshes with the new row; the new point also appears under **Smart Panel** in the navigation tree.

Field values may not contain the pipe character. If one does, the dialog shows: "Field "<label>" cannot contain a pipe character (|)." Other problems are reported inline in the dialog (or "Add failed." if no specific reason is available).

Editing a data point

Click a point's name in the **Data** column. The edit dialog opens, titled **Edit "<point name>"** — for example **Edit "Temperature 1"**. It offers the same per-point fields as the Add dialog: **Description**, **Modbus Register** and **Type**, with the live **Value** shown read-only. (The point's name appears in the dialog title rather than as a field.)

Click **Save** to apply your changes, or **Cancel** to close without saving. The table refreshes to show the updated point.

Deleting a data point

At the bottom left of the edit dialog is a red **Delete** button:

1. Click **Delete**. The edit dialog closes and a confirmation dialog titled **Delete Data Point** appears:
"Delete the data point "<name>"?" with the caution "This removes the mapping from the panel. The module input itself is unaffected."
2. Click the red **Delete** to confirm, or **Cancel** to keep the point.

On success the point disappears from the table and the tree. If the delete fails, the reason replaces the confirmation text (or "Delete failed." if no specific reason is available).

Note: Deleting a data point only removes the DM's mapping. Nothing changes on the CM06P module itself — you can re-add the same register later.

What View users see

View users get the same live page — the status badge, the warning banner, the color-coded table and the auto-refreshing values are identical. The differences are that the **Add** and **Properties** buttons do not appear, and point names in the **Data** column are plain text rather than links, so no editing dialogs can be opened.

Bringing Smart Panel inputs into units (the DIRS driver)

On an SPDM the wired Smart Panel — the Pythia CM06P I/O module — is the site's raw data source: dry contacts, run signals, temperatures and analog sensors come in over a direct serial link and land on the **Smart Panel** page as a list of live inputs. But an input on that page is not yet something you can monitor, alarm, trend or hand to an upstream integration; for that you need it to be a normal **Unit**, alongside every other piece of equipment on the site. The **DIRS** driver is the bridge that turns a Smart Panel input into a unit data point.

What the DIRS driver is

DIRS is an internal **source** driver. Unlike the SNMP, Modbus or BACnet drivers you met in Part 1, a DIRS unit talks to no hardware at all — it has no COM port and no IP address to poll. Instead, each of its data points carries a **Directory** setting that points at another location inside the DM, typically a Smart Panel input. On every poll the driver reads the live value at that location and copies it into the unit's own data point. The result is a unit whose points mirror the panel inputs you chose, and which then behaves like any other unit: it has a status, a dashboard tile, and points you can alarm and trend.

The workflow: mapping panel inputs into a unit

You already know how to **Add Unit** and add **Data Points** from Part 1. Mapping Smart Panel inputs is that same flow with two specific choices, done in a specific order.

1. **Add the input on the Smart Panel page first.** The input you are going to map has to exist as a source before you can point a unit at it. On the **Smart Panel** page, add the module input, set its register and type, and confirm a live **Value** appears. (This is the Smart Panel page covered just before this section — see "Adding a data point" there for the full procedure.) The inputs listed on that page are the source side of every mapping you are about to make.
1. **Add the DIRS unit.** Navigate to the group that should hold the unit and click **Add Unit**. Set the **Unit Name** (for example, "Room Contacts" or "AHU Manager"), choose the internal **Network** the DIRS driver runs on, and set **Driver** to **DIRS** — it appears in the dropdown as **DIRS (NET)**. (DIRS is a **NET**-type driver that talks to the loopback address `127.0.0.1` internally, not to any real device.) The **Network** and **Driver** dropdowns each show their entries as **Name (Type)**, and a short description line appears beneath the **Driver** dropdown as you select, so you can choose by purpose rather than by abbreviation. Click **Save**; the unit is created and the tree navigates to it, ready to configure.

> **Note:** DIRS is an internal driver, so it pairs with an internal (loopback) network, not a real field bus. When you save, the DM validates that the network and driver types are compatible; if you pick a mismatched network the save is rejected with a message naming the two types. Choose the internal network set up for this purpose.

1. **Add and bind each data point.** Open the new unit's **Data Points** page and click **Add**. Give the point a **Name**, a **Description** and a **Type** for how the value should be shown, then set its **Directory** to the panel input's path. You type this path in by hand, in the form `\Smart Panel1\<input>` — this is the binding that ties the unit's point back to the Smart Panel input. Make sure the point is **enabled**, then **Save**. From then on, every poll copies that input's current value into this point, and the value is available to the unit's status, alarms and trends like any other monitored point.
1. **Repeat for each input.** Add one data point per Smart Panel input you want to surface on the unit, then add alarms on the unit's **Alarms** page against those points as needed.

Caution: The order matters. If you skip step 1 and create the unit first, the **Directory** points at nothing and the data point simply shows no value. Add the input on the Smart Panel page before you bind a unit point to it.

Plain DIRS vs the DIRS_PYT_ management variants

The **Driver** dropdown offers plain **DIRS** and a family of pre-built **DIRS_PYT_...** variants. They solve two different problems.

- **Plain DIRS** is the general-purpose case above: an empty internal-source unit that you fill by hand, binding each data point to a Smart Panel input. Use it whenever you just need a handful of panel inputs to show up as a monitored unit.
- **DIRS_PYT_... variants** are pre-built cooling and exhaust-group managers. They already ship with their own data points, controls and alarms, so you do not hand-map anything. Instead of bringing in raw inputs, they coordinate a set of units you are **already** monitoring — handling lead/standby selection, failover, lead-lag rotation and temperature staging. You configure them through **properties** on the unit — chiefly which sibling units they manage, plus the lead/standby, rotation, failover and delay settings — rather than by mapping. The units a variant manages must be **siblings** in the same container.

When to use which: choose **plain DIRS** when you need panel inputs to appear as a unit; choose a **DIRS_PYT_...** variant when you are coordinating a group of cooling or exhaust units as a redundant team. If no fixed variant matches a site's unit count or logic, Pythia can supply a custom management driver.

Decoding the variant name: the suffix encodes the group shape. `AC` marks a cooling group and `EF` an exhaust-fan group; the first digit is the total units in the group and the second is how many normally run; `FA` means failover-capable; and a trailing digit is an internal behavior revision. So **DIRS_PYT_AC32FA2** reads as: air-conditioning group, three units total, two normally running, failover-capable, revision 2. The shipped variants:

Driver	Use it for	Units	Normally running
DIRS (plain)	Hand-mapping Smart Panel inputs into a unit	—	—
DIRS_PYT_AC2	Two cooling units, simple lead + backup	2	1
DIRS_PYT_AC21FA	Two cooling units with temperature staging	2	1
DIRS_PYT_AC21FA1	Two cooling units, alternate failover routine	2	1
DIRS_PYT_AC31FA / AC31FA2	Three cooling units, 1 running + 2 standby	3	1
DIRS_PYT_AC32FA / AC32FA2	Three cooling units, 2 running + 1 standby	3	2
DIRS_PYT_AC62FA4	Large six-unit cooling group with staging	6	up to 2
DIRS_PYT_EF2	Two exhaust fans, lead + backup	2	1

SPIGM

SPIGM runs the *identical software* to SPDM — the same console, the same Smart Panel workflow, and the same DIRS mapping. The only difference is the physical panel build (the hardware enclosure). Everything in this section applies to an SPIGM site exactly as written.

Part 4 — SPAC4 (Smart Panel AC-4 Controller)

SPAC4 is a Smart Panel that controls air-handling equipment. It replaces several core concepts — the section states what does and does not carry over from the Core Platform.

SPAC4 — Smart Panel AC-4 Controller

SPAC4 is a member of the DM family, but it is a different kind of product from the plain DM you met in Part 1. A plain DM *monitors* equipment: you build a tree of locations, groups, and units, map data points, and watch values. SPAC4 *controls* equipment. It drives a physical **Smart Panel AC4** — an AC, fan, and air-handler control panel — and runs the on/off staging logic itself, deciding when each unit turns on and off. This part documents only what is unique to SPAC4. Everything the shared platform already does works the same way here and is covered in Part 1.

A given DM is either a SPAC4 panel or a plain DM, never both. If you are administering a SPAC4 box, read this part alongside Part 1; if you are administering a plain DM, you can skip this part entirely.

How SPAC4 differs from the core platform

Most of what you learned in Part 1 applies unchanged. Two areas are replaced.

Note: SPAC4 keeps the shared platform for **alarms, automations, users, system settings, and integrations**. Those chapters in Part 1 describe SPAC4 behavior exactly — the alarm console, alarm categories and severities, automation rules, user roles and sign-in, the generic **Settings** pages, and the Modbus/TCP, SNMP, and BACnet/IP mappings all work here as documented there. Do not expect SPAC4 to change how any of those behave.

Two Part 1 chapters do **not** apply to SPAC4:

- **The Locations / Groups / Units hierarchy chapter.** A plain DM lets you build a `\Locations` tree of your own design. SPAC4 has no `\Locations` tree. In its place is a fixed, install-built topology (described next) that you cannot restructure — you only add and delete units within it.
- **The generic Properties-dialog workflow.** A plain DM edits container and unit settings through the standard property dialog, one dialog per item. SPAC4 edits its control settings inline on custom pages using a staged **Save / Cancel** model with server-side validation. The generic property dialog is replaced by that model on the SPAC4 pages (a small number of panel-wide and group-wide options still open their own **Properties** dialogs, but even those save the SPAC4 way).

Everything else in this part — the panel page, the group page, the unit page, the fan visual, the control logic, and the comms-loss behavior — has no equivalent in a plain DM.

The fixed Smart Panel topology

Where a plain DM shows `\Locations` at the top of the tree, a SPAC4 box shows a single **Smart Panel** node — the whole box. Beneath it the shape is fixed:

- **Smart Panel** — the top control node. There is exactly one, it is permanent, and it opens the **panel page**.

- **Group 1** and **Group 2** — always present, side by side. Both groups always exist; you cannot add a third group and you cannot remove either one. An administrator may rename a group, but that is all. Each group opens the **group page**.
- **Unit 1 ... Unit 4** — up to four units per group, in four fixed slots. Units are the only items you manage: you add them into empty slots and delete them from occupied ones. Each unit opens the **unit page**.

A **Settings** node holds the panel's fixed input/output wiring and protocol mappings — the panel inputs (such as *Unit N Fail*), the group sensors, the acknowledge button, the panel outputs and LEDs, and the Modbus/TCP, SNMP, and BACnet/IP mapping lists. These are pre-wired to match the hardware and a typical user never edits them. The **Settings** pages look and behave exactly like the core DM **Settings** pages from Part 1.

The two groups and all of the panel's control wiring are created automatically when the product is installed or updated. You never build this skeleton, and there is no way to build it wrong. Adding and deleting units is the only structural change you can make.

One more thing appears on every SPAC4 page: **a unit is drawn as a vented cabinet with a large fan**. The fan spins while the unit is running, turns red when the unit is in alarm, and greys out when the unit is off or disabled. This fan visual is your at-a-glance read of what the panel is doing.

The panel page

Opening the **Smart Panel** node shows the panel page: one **configuration console card per group**, side by side — a card for **Group 1** and a card for **Group 2**. This is primarily where you configure each group's control logic. The page refreshes live, about every two seconds, on its own; a status dot and an **"Updated: HH:MM:SS"** stamp at the top right confirm it is polling.

Each group card is laid out top to bottom:

- **Group title and status pill**. The title (for example **Group 1**) carries a status pill reading **Normal**, **ALARM**, or **Disabled**. Clicking the title — and only the title — opens that group's own page.
- **Units strip**. A row of small status tiles, one per unit assigned to the group. This row is hidden when the group has no units. Each tile shows a mini fan (spinning while the unit runs), the unit name, and the unit's mode, with a colored edge indicating its state. Clicking a tile opens that unit's page.
- **Live chips**. A read-only strip showing the group's current condition: the **Sensor** reading (shows -- when the group uses no sensor), the current **STAGE n**, a **▲ STAGING EXCEEDED** chip when the sensor has risen past the top setpoint, and the **Unit Hold** countdown (shows **Inactive** when nothing is being held).
- **Group Settings**. Every group setting, inline and editable, organized under four sub-headings: **Failover**, **Auto Sequencing**, **Staging**, and **General**. How you edit these is described under "Everyday use" below; what each one does is described next.

An administrator also sees a **Properties** button at the top right of the page, which opens the panel-wide **Smart Panel Properties** dialog (covered under "Everyday use"). The panel has no Delete — it is a permanent, single item.

Unlike a plain DM container page, which is just a list of child items, the panel page is a live control dashboard: group control logic surfaced as editable settings, live sensor / stage / hold chips, and fan tiles, none of which exist on the generic page.

The per-group control consoles

The four setting groups on each card are the heart of SPAC4. Here is what each control does in operator terms.

Staging brings units on and off as a measured value rises and falls — the core cooling logic.

- **Staging Sensor** selects what the group watches and whether it reads in Fahrenheit or Celsius. There are two choices: **10K F** and **10K C** — a 10K thermistor read in °F or in °C.
- **Staging Setpoint 1** through **Staging Setpoint 4** are the four thresholds. As the sensor climbs past each setpoint, another unit stages on; as it falls back, units stage off. The allowed range depends on the sensor's units: **45–95** when the **Staging Sensor** is a Fahrenheit sensor, and **7–35** when it is a Celsius sensor. Setting a value outside those bounds is rejected on save (see the validation examples below).
- **Staging Range** is the hysteresis band — the gap between the turn-on point and the turn-off point at a threshold. It stops a unit from rapidly cycling on and off when the sensor sits right at a setpoint.
- **Staging Alarm** raises an alarm when the group hits **STAGING EXCEEDED** (the sensor is above the top setpoint and the group has no more capacity to add).
- **Sensor Source** chooses where the reading comes from: **Local Sensor**, or **Use Other Group Sensor if Not Detected** (fall back to the other group's sensor when this group's own sensor is missing).
- The enable options for staging are **Disabled**, **Staging Up Enabled**, and **Staging Down Enabled**.

Failover provides lead/standby redundancy, so a single unit's fault does not take the whole group offline — a standby unit takes over for a failed primary. Two sub-settings control what happens after the primary recovers:

- **Latch Primary Alarm** — keep the failover latched (stay on the standby) after the primary's alarm clears, rather than switching back automatically.
- **Latch Primary On** — keep the primary running once it has been brought on, rather than dropping it back to standby.

Leaving both unlatched lets the group return to its normal lead as soon as the primary is healthy again; latching holds the failed-over state until you intervene, which avoids a second changeover disturbance.

Auto Sequencing rotates which unit leads the group, on a schedule, to even out run-hours and wear across the units.

- The rotation **interval** is chosen from a list ranging from hourly, through day-of-week and day-of-month options, to a running-days schedule.
- **Auto Sequencing Hour** and **Auto Sequencing Minute** set the time of day the changeover happens.
- **Auto Sequencing Rotate** rotates the lead immediately. This is a one-shot command, not a stored setting: you tick it and **Save**, the panel performs a single rotation, and the box clears itself.

General holds the group **Description** and any group settings not in the three control sections above.

The Unit Hold anti-short-cycle timer

Compressors and fans should not be switched on again the instant after they switch off — restarting a compressor too soon can damage it. SPAC4 enforces a minimum off-time with **Unit Hold**. When a unit turns off, it is *held* off for a set period before the panel is allowed to start it again. The panel page's live chips show this as a **Unit Hold** countdown, ticking down the remaining hold time; when nothing is being held it reads **Inactive**. The hold length is governed by the panel's **Unit Cycle Time** (set in **Smart Panel Properties**). You do not clear a hold manually — you watch it count down and know the panel is protecting the equipment.

The group page

Where the panel page holds each group's *settings*, the group's own page — reached by clicking a group title, or the **Group 1 / Group 2** node in the tree — is the group's **units console**. It is where you see and manage the units in that group.

- **Header.** The group name, its status pill, and an updated dot confirming the live refresh.
- **Live strip.** The same **Sensor / STAGE / STAGING EXCEEDED / Unit Hold** chips as the panel page, plus a **Failover** readout (Yes / No) and an **Auto Seq** readout. While auto sequencing is enabled it also shows **Last** (the last changeover time) and **Hours** (accumulated run-hours).
- **One card per assigned unit.** Each card carries the fan visual, the unit's run state — **RUNNING, OFF, ALARM**, or **DISABLED** — a status pill, and the unit's own settings inline: **Mode, Connection, Schedule, Schedule Hour, Schedule Run Time**, and **Description**. You edit these with the same staged **Save / Cancel** model used everywhere in SPAC4.

An administrator also has the unit-management controls here: **Add** to create a unit, **Delete** on each unit card, and a **Properties** button. Those are covered under "Everyday use" below.

If a group has no units, the console shows a message in place of the cards. An administrator sees: **"No units assigned to this group — use Add Unit to create one."** A View user sees the same message without the "use Add Unit" hint (they cannot add units).

The unit page

Clicking a unit — from a tile, a card, or the tree — opens the unit page, the detailed view of one fan or air-handler.

- **Header.** The unit name and its status pill, with a ← <Group> back-link to the group page and the unit's description shown as the subtitle.
- **Hero.** The unit drawn as a vented cabinet with a large fan. The fan spins while the unit runs, turns red on alarm, and greys out when the unit is off or disabled. Alongside it are the run state and two badges:
- **Mode** — one of **Primary**, **Standby**, **Standby Schedule** (shown as **STBY SCHED**), or **On**.
- **Connection** — the unit's connection state.
- **Unit Settings**, edited with the staged model (with **Save** and **Cancel** in the toolbar):
- **Mode** — the unit's role in the group: **Off**, **On**, **Standby**, or **Primary**.
- **Connection** — how the unit's control relay is wired: **Normally Open** or **Normally Closed**.
- **Schedule** — an optional scheduled run for a standby unit: **Disabled**, or a recurring schedule by **Day of Week** (Sunday–Saturday), **Day of Month** (1–31), or **Days** (a repeating day interval).
- **Schedule Hour** (0–23) and **Schedule Run Time** in minutes (0–59) — when a scheduled run starts and how long it runs.
- **Description** — free text.
- **Summary Alarm.** This is the unit's single alarm. A live state pill reads **ACTIVE**, **Normal**, or **Disabled**, and below it are the alarm's settings: **Enable**, **Category**, **Severity**, **Count**, and **Description**. The severity choices are **CRITICAL**, **MAJOR**, **MINOR**, **WARNING**, and **INFORMATION**. The alarm itself, and the alarm grid beneath the unit, are the standard DM alarm pages from Part 1.
- **Data Point — Unit On.** This is the unit's single data point: the live run value, plus its settings **Enable**, **Trend**, and **Description**. Turning on **Trend** feeds the standard DM trend and history, so a unit's on/off record is trendable like any other DM data point.

An administrator sees a **Delete Unit** action in the toolbar (covered below). The data point and alarm grids under the unit are the standard DM pages.

The fan hero, the live run state, the Mode and Connection badges, and the fixed single-Summary-alarm-plus-single-Unit-On layout are all SPAC4-specific — a plain DM unit page has none of them.

About every two seconds — and NO COMMS

All three SPAC4 pages poll on their own, refreshing roughly every two seconds, so values, fans, chips, and countdowns update without you doing anything.

Caution: When the DM loses communication with the panel, every SPAC4 page **dims, freezes its animation, and shows a red "NO COMMS" badge**. A frozen fan with the **NO COMMS** badge means the values on screen are stale — the last known state, not the live one. Do not act on the displayed readings until the badge clears and the page resumes refreshing.

Everyday use — adding units, editing, and saving

The staged Save / Cancel model

Every SPAC4 settings field — on the panel page, the group page, and the unit page — is edited the same way, and it is different from the core DM property dialog.

1. **Change a field.** Type a number, pick from a list, or tick a box. The field is now **pending**: a small amber dot appears beside it, and the card's or section's **Save** and **Cancel** buttons light up. Nothing has been written yet.
2. **Read the allowed range.** Each field shows its allowed range beside the label, taken from the panel's schema — for example a minute field hints "**0–59 min**". Some ranges depend on another field: changing a list field can change what the others allow. Setting the **Staging Sensor** to a Fahrenheit or Celsius sensor, for instance, re-hints the **Staging Setpoint** limits to **45–95** or **7–35** accordingly, and the grid refreshes those hints after you make the change.
3. **Click Save.** Validation runs on the panel. On success a green "**Saved ✓**" appears briefly and the grid re-reads the panel's stored values.
4. **Or click Cancel.** This discards every pending edit and snaps the fields back to their stored values.

When validation fails, the panel reports the reason **in place at the offending field** — a red **X** next to the field, with the panel's message. The messages, each prefixed with the field's own name, read verbatim:

- "***...must be a number***"
- "***...is below the minimum of <n> ***"
- "***...is above the maximum of <n> ***"
- "***...value "<x>" is not one of the allowed options***"
- "***<field> cannot contain the | character.***"

So typing 99 into a Fahrenheit setpoint produces, in place, "**Staging Setpoint 1 is above the maximum of 95**". Fix the flagged field and click **Save** again; other pending edits stay pending until the whole save succeeds.

Adding a unit

Admin only: Adding units requires Administrator access.

On a group page, click **Add** to open the **Add Unit** dialog. Pick a unit position and fill in the new unit's settings, then confirm. Two rules govern the position picker:

- **Unit names are unique across the whole panel — both groups.** The picker offers only the positions still free across both groups combined.
- When every slot is taken, the group page shows: **"*All four unit positions are assigned (units are unique across both groups).***"

If an add fails because the name is already in use elsewhere, you see: **"*Add failed (is the unit name unique across both groups?).***"

Deleting a unit

Admin only: Deleting units requires Administrator access.

You can delete a unit two ways: the **Delete** control on the unit's card on the group page, or **Delete Unit** in the unit page toolbar. Either way the panel asks you to confirm:

Caution: Deleting a unit removes the unit and **all of its alarms and data**. The confirmation reads: **"*Delete <Unit> and all of its alarms and data?***" After you confirm from the unit page, the DM returns you to the group page. This cannot be undone.

The panel Properties

Admin only: The panel **Properties** dialog is available to Administrators only.

The **Properties** button on the panel page opens **Smart Panel Properties** — the panel-wide options that apply across both groups:

- **Enable/Silence Buzzer** — turns the panel's audible buzzer on (**Enable**) or silences it (**Disabled**).
- **Acknowledge Alarm and Unlatch** — a toggle that acknowledges the panel's current alarm and clears any latched state.
- **Alarm Delay Time in Seconds** — how long a condition must persist before the panel raises the alarm.
- **Unit Cycle Hold Time in Minutes** — the minimum off-time that drives the **Unit Hold** anti-short-cycle countdown.
- **Panel Status** — a read-only indicator of the panel's current state (for example **Normal**).

Edit the fields and click **Save** (or **Close** to dismiss without saving). There is no Delete — the panel is permanent.

Restore Defaults

Admin only: **Restore Defaults** lives in the group **Properties** dialog and is available to Administrators only.

Each group has a **Properties** dialog, reached from the group's own page, that edits the same group settings shown inline on the panel page — so the group control settings are reachable both ways. This dialog also holds the one destructive group action, **Restore Defaults** (tooltip: "*Reset every group setting to defaults and remove all units from this group*").

Caution: **Restore Defaults** resets **every** setting in the group to its defaults **and removes all units** assigned to the group. It cannot be undone. The confirmation reads: "*Restore <Group> to its default configuration? This resets EVERY group setting to defaults and REMOVES ALL UNITS assigned to this group.*"

On success the group's settings return to their defaults, its units disappear, and the tree updates. On failure you see: "*Restore failed: <reason> *".

Administrators versus View users

The live views are identical for everyone. Group consoles, unit cards, the fan hero, the live chips, and every auto-refreshing value look the same and update the same for a View user as for an administrator — a View user sees the full live state of the panel at all times.

The difference is control. A View user can change nothing:

- **Every settings field is shown disabled.** Values are visible but cannot be edited or saved.
- The **Add**, **Delete** / **Delete Unit**, **Restore Defaults**, and **Properties** buttons are **hidden**.
- Because there is nothing to stage, the **Save** and **Cancel** buttons never appear for a View user.

An administrator gets all of those controls, plus the **Save** and **Cancel** buttons on each settings card and section.

Note: SPAC4 also offers an anonymous, full-screen **panel display** — a read-only "kiosk" face for the panel's own tablet or any browser on the panel's network, needing no sign-in. It shows the whole panel at a glance: both groups side by side, each unit as a fan tile with its role (**PRIMARY** / **STANDBY**) and state (**RUNNING** / **OFF**), the group's Failover, Auto Sequencing, Last Changeover, Run Hours, Unit Hold and Staging readouts, and a bottom status strip of **ALARM**, **ACK**, **POWER** and **COMMS** lamps. A **Console** button in the corner opens the signed-in DM. This display is read-only and has no controls — it is a separate surface from the signed-in console pages above.

Appendices

Appendix A — Message Reference

This appendix collects the error and confirmation messages you are most likely to meet in the DM console, grouped by the area of the console where they appear. It is a lookup aid, not a complete catalog — routine field prompts and messages that explain themselves are omitted. Placeholders such as `<name>`, `<ip>`, and `N` stand for the actual name, address, or count shown on screen.

Two message styles exist side by side. Older messages embed a four-digit code in the text itself ("0003 - Access denied"); newer messages show plain text with a separate five-digit code (for example, code 20122). When you contact support, quote the code if one is shown.

Note: "0003 - Access denied" (or "Access denied" with a code) anywhere in the console means the same thing: you are signed in as a View user, or your session has expired. Sign in as an Administrator and retry.

Signing in and access

Message	When it appears	What to do
"Invalid credentials" (20122)	The sign-in attempt failed. The console deliberately shows the same message for a wrong email, a wrong password, and a disabled account.	Check your email and password and try again. An Administrator can see the specific reason in the Logons log.
"Nonce expired" (20123)	You clicked Sign In more than 60 seconds after the sign-in page prepared your attempt.	Simply try signing in again — the page issues a fresh challenge.
"Not logged in"	You opened a page that requires a signed-in session (for example the Users list) after your session expired.	Sign in again.
"0003 - Access denied"	A View user (or an expired session) attempted a change — add, rename, delete, clear, save, or toggle.	Sign in as an Administrator.
"0001 - IP Address not allowed: <ip> "	Your computer's IP address is not on the DM's access list.	Ask an Administrator to add your address in the console's IP access settings.
"0002 - Unknown command: <cmd> "	The page and the DM software are out of step, usually right after a software update.	Hard-refresh the browser (Ctrl+Shift+R). If it persists, contact your Administrator.
"0501 - COM object creation failed"	The DM's web components are not registered — a deployment problem, not a user error.	Contact your Administrator (the server-side fix is re-registering the DM's web component).
"Server secret not provisioned" (20125)	Sign-in or password change fails because the installer did not complete its security setup.	An Administrator should re-run the DM installer.
"Crypto error" (20126)	An internal security calculation failed on the server.	Retry once; if it persists, report it to your Administrator.

Locations, Groups, and Units

Message	When it appears	What to do	
"Location name is required." (also "Group name is required.", "Unit name is required.")	You clicked Save in an Add modal with the name blank.	Enter a name.	
"Name must be 100 characters or less." (server: "Location name must be 1-100 characters", 10001)	The name is too long (or empty on the server side).	Shorten the name to 100 characters or fewer.	
"Name cannot contain: \ & % # ? " , " (10002)	The name includes a character the DM does not allow in container names.	Remove the listed characters and save again.	
"A location with that name already exists" (10003)	Adding or renaming a location to a name already in use.	Pick a different name.	
"A group with that name already exists under this location" (10003)	Adding or renaming a group to a sibling's name.	Pick a different name — group names only need to be unique within their location.	
"A unit with that name already exists in this group"	Adding a unit with a name already used in that group.	Pick a different name.	
"Unit name cannot contain backslash, slash, or pipe."	The unit name includes \ , / , or \	`.	Remove those characters.
"Please select a network." / "Please select a driver."	You clicked Save in Add New Unit without choosing both dropdowns.	Choose a network and a driver, then save.	
"Driver type (<type>) does not match Network type (<type>)"	The driver you picked speaks a different connection type than the network you picked (for example a LAN driver on a serial port). This is checked when you save, not while you pick.	Choose a driver and network of the same connection type (network drivers on LAN networks, serial drivers on COM ports).	
"Location contains N group(s). Remove all groups before deleting." (20003) — shown as "The location "	You tried to delete a location that still has groups in it. The delete is refused.	Delete (or move the contents of) each group first, then delete the location.	

Message	When it appears	What to do
<name> " contains N group(s)."		
"Group contains N item(s). Remove all items before deleting." (20009) — shown as "The group " <name> " contains N unit(s)."	You tried to delete a group that still has units.	Delete the units first, then the group.
"Location does not exist" / "Group does not exist under this location"	The thing the page is showing was deleted — usually by another user — between refreshes.	Nothing; the page redirects itself (see next row).
"Redirecting to Locations in 3 seconds... Go now"	Follows the message above: the console counts down and returns you to the nearest surviving level.	Wait for the redirect or click Go now .
"Unable to load location data. Check the API connection."	The Site Overview page could not reach the DM twice in a row on a cold load.	Check that the DM service is running and your network connection to it; then refresh.
"Saving..." (full-screen overlay)	After any add, rename, or delete — the console is waiting for the server to finish serving the changed item. Normally under a second.	Wait; the console navigates automatically when it clears.
"Are you sure you want to delete the location " <name> "?" — with "This action cannot be undone."	Every delete asks for confirmation with this pattern (location, group, unit, and similar).	Click Delete only if you are certain — there is no undo.

Properties and renaming

Message	When it appears	What to do		
"Property " <name> " cannot contain the \"	character."	A value you typed in a Properties modal contains a pipe character, which the DM reserves.	Remove the \"	` and save again.
"Could not load properties."	The Properties modal could not fetch the item's fields.	Close and reopen the modal; refresh the page if it persists.		
"Save failed."	A property save failed without a more specific server message.	Retry; if it persists, report it to your Administrator.		
"This container has no editable properties."	You opened Properties on an item with nothing to edit (common on the site root).	Nothing — informational only.		
"Could not open parent for rename" (10004)	A rename hit a transient server problem.	Retry; report it if it keeps happening.		
"Rename failed" (10005)	The rename could not be completed on the server.	Retry; if it persists, report it to your Administrator.		
"Source location does not exist" (10006, wording varies by item)	The item you are renaming was deleted meanwhile (usually by another user).	Refresh the page.		

Alarms

Message	When it appears	What to do
"0101 - Database connection failed" / "0102 - Database query failed"	Alarm History could not be read — the DM's SQL database is down or unreachable.	An Administrator should check that the database service is running. Active and Information alarms are unaffected.
"1401 - Delete failed"	Clearing alarm history failed at the database.	An Administrator should check the database service, then retry.

Automations

Name and address checks on automation screens use the same rules as elsewhere: names must be 1–100 characters (10001), must not contain the invalid characters listed above (10002), and must not

duplicate an existing sibling (10003). The message text names the field, for example "Automation name must be 1-100 characters".

Message	When it appears	What to do
"Type must be Alarm or Heartbeat" (20038)	Creating an automation without a valid type selection.	Choose Alarm or Heartbeat in the Add dialog.
"Email address is required" (20046)	Adding an email recipient with the address blank.	Enter the recipient's address.
"Must be a valid email (e.g. user@example.com)."	The address you typed is not a valid email format.	Enter a full address, such as <code>user@example.com</code> .
"IP address is required" (20053)	Adding an SNMP trap destination with no address.	Enter the destination IP.
"IP address must be a valid IPv4 address (e.g. 192.168.1.10)" (20054)	The trap destination is not a well-formed IPv4 address.	Enter a dotted address such as <code>192.0.2.10</code> .
"Day of week is required" (20060)	Adding a time exception without choosing a day.	Pick a day of the week.
"Start time is required" / "Stop time is required" (20069 / 20070)	A time-exception range is missing its start or stop.	Enter both times (24-hour HH:MM).
"Stop time must be after start time" (20071)	The stop time is at or before the start time.	Enter a range where the stop is later than the start.
"This time range overlaps the existing range 08:00 - 12:00" (20072)	The new range overlaps one already defined for that day.	Adjust the times so the ranges do not overlap, or edit the existing range.

SMTP Relay (email delivery)

Message	When it appears	What to do
"SMTP Relay service not installed" (20480)	Opening Settings → SMTP Relay on a gateway where the relay service is not running — typically an older gateway that predates it.	The relay is enabled by default on current systems; if you see this, contact Pythia Technologies support.
"Access denied" (20481 / 20483 / 20486 / 20487)	A View user tried to save relay settings, send a test message, reset counters, or edit the allow list.	These actions require Administrator access.
"<key> value is invalid" (20482)	A relay configuration value — a number, a port, or a listen-port — is out of range or malformed when saving.	Correct that field; a blank field means "use the default."
"A valid destination address is required" (20484)	Sending a test message with the address blank.	Enter an address to send the test to.
"Relay refused the message: <error> " (20485)	The upstream mail server rejected the test message.	Read the server's reason, then check the Smart Host , Port , and allow list.
"An entry (IP, range, or hostname) is required" (20488)	Adding an allow-list entry with the field blank.	Enter an IP address, CIDR range, or hostname.
"Not a valid IPv4 address, CIDR range, or hostname" (20489)	The allow-list entry is not one of those forms.	Use an address like <code>192.0.2.10</code> , a range like <code>192.0.2.0/24</code> , or a hostname.

Trends

Message	When it appears	What to do
"✓ Saved"	Flashes next to a Trend bucket dropdown in Configure Trended Points after a per-point change is saved.	Nothing — confirmation. The trend service creates or removes the trend shortly after.
"Could not update trend."	A per-point bucket change failed to save.	Retry; if a specific server message is shown instead, follow it.
"Turn OFF trending for N points?"	Bulk-applying Off in Configure Trended Points — always asks first.	Confirm to stop trending the selected points (their history is kept), or cancel.
"Set Trend " <bucket> " on N points?"	Bulk-assigning a bucket to more than 20 points at once.	Confirm or cancel.
"Saved N points." / "Saved X, failed Y (<first error>). Failed rows stay selected."	After a bulk apply finishes.	On partial failure the failed rows remain checked — fix the cause and click Apply again.
"No trend data in this window"	The chart has no samples in the selected time window — normal on a brand-new trend that hasn't collected yet.	Wait for the service to collect samples, or widen the Window .
"Trend leaf is not GUID-stamped" (20100)	A trend is missing its internal identity — a post-install condition that should not occur in normal use.	Report it to your Administrator (re-running the installer heals identities).
"Database connection failed" (20101) / "Database query failed" (20102)	Trend history could not be read from the database.	An Administrator should check the database service.
"Delete failed: <server message> "	Deleting a trend failed.	Read the embedded message; retry or report it.
"All historical data values collected for this trend will be permanently deleted. ... This action cannot be undone."	The Delete Trend confirmation.	Confirm only if you accept losing the history — re-enabling the point later starts a fresh, empty trend.
"Could not rasterize chart for export." / "PDF library failed to load (jsPDF)."	A PNG or PDF export failed in the browser.	Retry; hard-refresh the page if it persists. CSV/Excel exports are unaffected.

Backups and Updates

Message	When it appears	What to do
"403 Admin access required"	A View user tried to download a backup file or upload an update file.	Sign in as an Administrator — these two file operations are Admin-only.
"400 Invalid filename"	The requested backup filename contains characters or path elements the DM refuses.	Use the download links in the Backups list rather than typing a filename.
"Only .bk files allowed"	A backup download was requested for a file that is not a .bk backup file.	Download only files shown in the Backups list.
"404 Backup file not found"	The backup file no longer exists — usually deleted since the list was loaded.	Refresh the Backups list.

Logs

Message	When it appears	What to do
"1002 - Database connection failed" / "1003 - Database query failed"	The log viewer could not read entries from the database.	An Administrator should check the database service.
"1102 - Delete failed"	Clearing a log's entries failed.	Retry; if it persists, check the database service.

Network

Message	When it appears	What to do
Invalid IP/netmask/gateway (20397)	Saving a network interface with a malformed address, netmask, or gateway.	Correct the address fields (dotted IPv4, e.g. 192.0.2.10) and save again.
Invalid DNS address (20398)	A DNS server entry on the interface is not a valid address.	Correct the DNS entry and save again.

Caution: Apply & Restart on the network settings page is not an error condition, but it reboots the DM server — the console drops for all users until the server is back up, and the browser is redirected to the new address.

Appendix B — Access Rights

The DM has exactly two access levels, set per account: **Administrator** (shown as **Admin**) and **View**. View accounts are read-only — they can browse, search, filter, switch between Cards and Table, open Properties modals, and export, but every change (add, rename, delete, clear, toggle, save, command) is refused with an Access-denied message. Administrator accounts can do everything. The console hides most Admin-only controls from View users (Add buttons, rename pencils, Save and Delete buttons), but the restriction is enforced by the server either way — and an expired session is treated exactly like a View user, so if a change is unexpectedly refused, sign in again. The sign-in page itself is what protects viewing: nothing in the console is visible without an account.

What each level can do, by area

Area	View users can	Administrators additionally can
Monitoring (Site Overview, Locations, Groups, Units lists)	Browse every level, search and filter lists, switch Cards/Table views, refresh, open Properties modals read-only	Add, rename, and delete locations, groups, and units; edit their properties (Enable, Code, Description, the Header 1–5 custom columns)
Units (unit detail)	View live data, data-point charts and exports, alarm status and history, and the controls list	Send control commands (On/Off, setpoints, Override, Clear Setpoints); add and edit data points, controls, and alarm settings; use the bulk Data Point Properties and Configure Alarms grids; clear alarm history
Dashboards (equipment driver pages)	View the live diagrams, status pills, and fleet views; open live chart modals	Edit the underlying unit and device properties
Alarms	Browse Active, Information, and History; filter by severity; search; open details; export	Clear individual or selected alarms, clear entire buckets, and trim or clear the alarm history
Automations	View automations, email recipients, trap destinations, time exceptions, and heartbeat schedules	Create and rename automations; add, edit, and delete recipients, trap destinations, time exceptions and ranges, and heartbeat intervals/periods
Trends	Browse buckets and trend charts; change window, chart type, and theme (personal, browser-local settings); export CSV/Excel/PNG/PDF; open Configure Trended Points read-only	Assign points to buckets or turn trending off (per point or in bulk); delete a trend together with all of its history
Smart Panel	Watch the live input monitor	Change the panel configuration (buzzer, delay, category, network); add, edit, and delete input points
Users	Manage their own account only — profile, password, sign-in email; the Users list shows only their own row	See and manage all accounts: add users, enable/disable, set access levels, reset passwords, rename, and delete (subject to the Admin-account protections below)
Settings	View the settings pages, license/serial, the backups list, and the system logs	Run Backup Now , download and delete backup files, upload and apply <code>.upd</code> software updates, set date and time, edit password policies, enable/disable/clear logs, and reboot the DM

Area	View users can	Administrators additionally can
Integrations	View the SNMP/Modbus/BACnet mappings and drill-downs, the ALC/WebCTRL URL pages, and all exports (CSV, PDF, Excel, HTML, JSON, MIB, EDE)	Set the Modbus register format and SNMP value format, use Enable All R/W , and reset the DM mapping IDs

Note: exports are generally available to View users, with one exception — downloading backup files (and uploading update files) is Admin-only and returns "403 Admin access required" otherwise. Also note that the server filters the Users list itself: a View user's browser never even receives other users' names.

The Admin account's protections

The built-in **Admin** account is the recovery account, and the DM enforces a set of invariants around it:

- The Admin account **cannot be renamed** — not from its own My Account page, and not by another administrator.
- The Admin account **cannot be deleted**.
- The Admin account **cannot be edited** through the Users page; the console renders its fields read-only.
- Administrators manage **their own** account only through **My Account** (self-service). The Users-page actions (rename, delete, reset password) refuse to act on the account you are signed in with.
- **Passwords are never displayed or exported** for any account — the DM stores only password hashes, and no page or export can retrieve a password.

Appendix C — Glossary

Terms you will meet throughout the DM console and this guide, in alphabetical order. Where a term has a different meaning in two areas (for example *bucket*), both are listed.

Access level — One of the two permission levels assigned to every account: **Admin** (full control) or **View** (read-only). See Appendix B for the full breakdown.

Alarm event — A single occurrence of an alarm condition, listed in the Active, Information, or History buckets with its timestamp and details. Events are what you clear; they are produced by alarm settings.

Alarm setting — The definition configured under a unit that tells the DM when to raise an alarm on a point (condition, severity, category). One alarm setting can produce many alarm events over time.

ALC/WebCTRL URL — A console page that lists ready-made URLs for pulling DM points into an ALC WebCTRL building-automation system.

Automation — A rule that sends notifications. An **Alarm** automation reacts to alarm events with emails and/or SNMP traps; a **Heartbeat** automation sends scheduled proof-of-life emails.

.bk file — A DM configuration backup file, created by **Backup Now** or the backup schedule and listed under Settings > Backups. Administrators can download and delete them.

Breadcrumb — The path line in a page's subtitle showing where you are in the hierarchy, e.g. **Sunbury Office > Environmentals**.

Bucket (alarms) — One of the three alarm lists: **Active** (current alarms), **Information** (informational events), and **History** (the searchable archive).

Bucket (trends) — One of the five fixed trend containers, each with its own sampling interval: 5 minutes, 15 minutes, 30 minutes, 1 hour, or 1 day. A data point is trended into one bucket at a time.

Category — A classification label on an alarm setting, used to group related alarms (visible in the Configure Alarms grid).

Console — The DM's web interface — the tree, top bar, status bar, and content pages this guide describes.

Control — A writable point on a unit through which you send commands: On/Off states or numeric setpoints. Control commands take effect as soon as they are sent, without a confirmation step.

Data point — A single value the DM reads from a unit, such as a temperature, a voltage, or a status. Data points appear in the unit's Live Data list and can be alarmed and trended.

DM — The Pythia Technologies Data Manager: the monitoring and control gateway this guide covers. It polls your equipment, raises alarms, records trends, and republishes data over SNMP, Modbus, and BACnet.

Driver — The software module that speaks a device's protocol. You pick a driver (and a matching network) when adding a unit; the driver determines which data points, controls, and alarms the unit offers.

Fail count — On an alarm setting, the number of consecutive polls a condition must persist before the alarm is raised — a guard against one-off glitches.

Gateway — Used two ways: the DM itself acts as a protocol gateway (republishing monitored data to external systems), and in network settings **gateway** is the router address an interface uses.

Group — The middle level of the hierarchy: an area or category within a location, such as **Environmentals** or **Power Units**. Groups contain units.

Heartbeat — A scheduled proof-of-life email sent by a Heartbeat automation, so recipients know the DM (and its site connection) is alive. Sent at fixed intervals or at set times of day.

Instance (BACnet) — The numeric object identifier assigned to a DM point in the BACnet mapping; external BACnet systems address the point by this number.

Interval — In a heartbeat automation, a fixed repeat spacing ("every N hours") at which heartbeat emails go out. Compare **Period**.

Location — The top level of the hierarchy: a physical site or building, such as **Sunbury Office**. Locations contain groups.

Mapping — The per-protocol table (SNMP, Modbus/TCP, or BACnet/IP) that exposes DM points to external systems, giving each point an ID, address, and format.

Network — A configured communication path — a LAN connection or a serial port — that units attach to. A unit's driver and network must be of the same connection type.

OID — Object Identifier: the dotted numeric address (e.g. 1.3.6.1.4.1. . .) by which an SNMP system reads a DM point from the SNMP mapping.

Period — In a heartbeat automation, a set time of day (HH:MM) at which a heartbeat email is sent daily. Compare **Interval**.

Register — The Modbus register address assigned to a DM point in the Modbus/TCP mapping; external Modbus systems poll the point at this address.

Setpoint — A numeric target value sent to equipment through a control — for example a temperature setpoint on a cooling unit.

Severity — The importance level assigned to an alarm, shown with a matching color throughout the alarm pages and used for filtering.

Smart Panel — The CM06P input-monitoring panel; its console page shows the panel's live inputs and (for Administrators) its configuration.

Status rollup — How alarm status climbs the tree: when any unit is in alarm, its group and location show **Alarm** too, so trouble anywhere lights up the whole chain. Disabled items are excluded from the rollup.

Time exception — A day-of-week plus time-range entry on an automation defining when it should *not* send notifications. Ranges on the same day may not overlap.

Trap — An SNMP notification message sent by an automation to a configured destination IP address when an alarm event occurs.

Tree — The navigation pane on the left of the console, mirroring the Locations > Groups > Units hierarchy plus the Alarms, Automations, Trends, Views, Users, and Settings branches.

Trend / Trend leaf — The stored history for one data point in one trend bucket. History is keyed by an internal GUID, so it survives renaming the location, group, or unit; deleting a trend permanently removes its history.

Unit — The bottom level of the hierarchy: a single monitored device, such as **Office TH1**, **CRAC-01**, or **UPS-21A**. Units carry the data points, controls, and alarm settings.

.upd file — A DM software update package. An Administrator uploads it under Settings > Updates and applies it; applied updates cannot be rolled back.

Window (trends) — The span of history a trend chart displays (24–168 hours or 30–180 days, depending on the bucket). It is a personal viewing preference stored in your browser — it does not affect what the DM collects or keeps.