

InfluxDB 3 Core with Command Center and Explorer

Powered by Code Creator | Complete Customer Installation and Operating Manual

A practical, page-by-page guide to launching the appliance, reading your dashboard, exploring telemetry, connecting other servers, running diagnostics, and protecting data with optional S3 backup and Guided Recovery.

Start here You can use Command Center, InfluxDB 3 Explorer and built-in EC2 telemetry without attaching an IAM role. An EC2 IAM role and a private S3 bucket are needed only when you choose to enable the optional backup and recovery workflow.

Guide at a glance

Chapter	What you will learn
1. Launch and sign in	EC2 settings, ports, first boot, password, HTTPS and URLs
2. Command Center home	Every Quick Action, status card, connection summary and Sign out
3. InfluxDB 3 Explorer	Default server and databases, running SQL queries, visualizing data
4. Add Telemetry Source	How optional remote Linux/EC2 sources differ from built-in metrics
5. System Health Check	Run Health Check and interpret PASS, ATTENTION, FAIL and STAGED
6. AWS Configuration Settings	Private S3 bucket, IAM role and copyable policy templates
7. Backup & Recovery	Configure Backup, schedules, Save & Test and Run Backup Now
8. Guided Recovery	Select backup, check safety gates, begin and verify a restore
9. Administration and help	Basic commands, common problems and security reminders

Important URLs in this guide use YOUR_PUBLIC_IP as a placeholder. Substitute the public IPv4 address assigned to your own EC2 instance. Do not use an IP address shown in a Code Creator screenshot.

1 Launch and sign in

1.1 Tested deployment configuration

Setting	Customer instruction
Operating system	Ubuntu 26.04.1 LTS; x86_64 architecture.
Tested EC2 instance	t3.xlarge (16 GiB memory class). Smaller sizes were not established as supported by this release test.
Tested root EBS volume	30 GiB. Increase storage to fit your telemetry ingestion and retention requirements.
Network and identity	Select an EC2 key pair, enable a reachable public IPv4 address for the public-IP quick start, and use a security group with the necessary inbound ports.
IAM role on launch	Not required for Command Center, Explorer or built-in telemetry. Optional when activating S3 backup and recovery.

1.2 Security-group inbound ports

TCP port	Purpose / recommended access
22	SSH administration. Restrict source to your public IP or trusted administrator CIDR.
80	HTTP used by the appliance's IP-certificate issuance/renewal process; allow public validation traffic when required for certificate management.
443	Command Center HTTPS. Restrict to trusted user networks where practical.
8443	Explorer HTTPS through the authenticated front door. Restrict to trusted user networks where practical.
8181	Internal InfluxDB engine; do NOT make this a public security-group inbound rule. The host binding is 127.0.0.1.

Remote collection If you add a remote telemetry source, follow that source page's connection instructions before opening any additional telemetry-ingestion ports. Do not expose internal management ports indiscriminately.

1.3 First-boot login

1. Launch the AMI, select your EC2 key pair, and allow first-boot initialization to finish. The first boot creates appliance-specific credentials and configures the services.
2. From your computer, connect by SSH using the Amazon private key selected during launch and the user name **ubuntu**. On Windows, PuTTY can be used for SSH.
3. On the EC2 instance, display the generated login file:
`sudo cat /root/FIRST_LOGIN.txt`
4. Open Command Center in your browser using `https://YOUR_PUBLIC_IP/` and sign in with the username and password shown in the file. The default username is admin; use your own generated password.
5. From Command Center choose Open InfluxDB 3 Explorer; it opens at `https://YOUR_PUBLIC_IP:8443/` using the existing authenticated gateway, without requiring a second Code Creator login.

Protect credentials FIRST_LOGIN.txt contains sensitive login information. Do not email or post it in tickets, screenshots or public chat. The password generated on your instance differs from demonstration credentials.

1.4 HTTPS and address changes

First boot attempts a trusted public-IP HTTPS certificate and can fall back to a self-signed certificate if issuance is unavailable. Check the HTTPS certificate line in FIRST_LOGIN.txt. If your public IP changes, the original IP-specific certificate and access URL may no longer be valid; consult your administrator before relying on an old link. Port 80 may be needed for certificate issuance and renewal.

2 Command Center home page: every control

Command Center is the primary dashboard. It summarizes database health, telemetry freshness and EC2 resource use, and offers shortcuts to the appliance's supported workflows. Its status display auto-refreshes approximately every 15 seconds.

2.1 Top-of-page indicators

Home-page element	Meaning / what to do
InfluxDB Healthy	The database responds to the dashboard's health check. If the status becomes unhealthy, open Run Health Check.
Telemetry state: RECEIVING	Recent measurements are entering the database. Fresh installations collect the appliance EC2 host automatically.
Latest sample: N seconds ago	How long ago the latest telemetry sample arrived. The age should advance and reset as new measurements are received.
Database telemetry	The default private InfluxDB database used for built-in host metrics.
Auto-refresh 15s	The dashboard updates its status automatically; you normally do not need to reload the page.

2.2 Quick Actions: what each button does

Button / link	What opens and when to use it
Open InfluxDB 3 Explorer	Opens https://YOUR_PUBLIC_IP:8443/ to view databases, run SQL queries and explore measurements. The default server and telemetry database should be ready automatically.
Add Telemetry Source	Opens the guided page for connecting another supported Linux/EC2 server and checking that it is reporting. It is not needed for built-in EC2 host telemetry.
Run Health Check	Opens System Health Check. Click its Run Health Check button to execute live database/data-flow and appliance-service diagnostics.
Backup & Recovery	Shows saved destination, schedule, most recent backup, Run Backup Now, Start Guided Recovery and Recovery Verification.
AWS Configuration Settings	Opens the optional backup-authorization helper: how to create or choose a private S3 bucket and attach an EC2 IAM role scoped to the bucket.

2.3 Operational Overview cards

Card	Meaning
InfluxDB Health — HEALTHY	The appliance's database is responding; this does not guarantee every external data source is reporting.
InfluxDB — Core	Identifies the included single-node InfluxDB 3 Core engine.
Telemetry — RECEIVING	New measurements are arriving in the appliance's telemetry stream.
Telemetry Database — telemetry	The named database receiving built-in EC2 host measurements.
CPU Active — percentage	Current CPU utilization reported for the appliance EC2 host; it is not a fleet-wide average.
Memory Used — percentage	Current memory utilization reported for the appliance EC2 host; it is not all connected servers combined.

2.4 Connection Summary and Live Status Notes

Connection Summary displays the internal InfluxDB URL (<http://influxdb3-core:8181>), the telemetry database and the private telemetry appliance mode. That hostname is internal to the Docker network. It is NOT the customer browser URL: keep port 8181 closed to the internet.

Live Status Notes restate whether InfluxDB is responding, telemetry is fresh, how old the latest sample is, and that the dashboard is authenticated. Use them as at-a-glance status, not as a substitute for a detailed diagnostic.

2.5 Sign out

Choose Sign out when finished, especially on a shared computer. Return to https://YOUR_PUBLIC_IP/ and sign in again when needed.

3 InfluxDB 3 Explorer: data and queries

Open Explorer from the Command Center Quick Actions. Code Creator configures the default server during first boot, so the normal customer path does not require manually adding a server, supplying another token, or creating the telemetry database.

3.1 Databases page

In Explorer, open the Databases / Manage Databases area to inspect available databases. The fresh validated appliance displayed telemetry and `_internal`. The telemetry database contains the bundled Telegraf system metrics; `_internal` is the internal InfluxDB database. Do not delete either during initial setup.

Explorer may offer database-creation, deletion, retention or other management controls. Those operations change database state. Use them only when you intend to make that specific change, and maintain a backup before destructive operations.

3.2 Run your first query

1. Open Explorer through the Command Center and navigate to its SQL query/editor area. Select or confirm the default telemetry database when prompted.
2. Enter the SQL below into the editor, then use the query execution control to run it.

```
SELECT time, host, usage_active
FROM cpu
WHERE time >= now() - interval '2 minutes'
AND cpu = 'cpu-total'
ORDER BY time DESC
LIMIT 10;
```

3. Review the returned time, host and usage_active columns. Recent timestamps confirm the appliance has collected CPU measurements. Query times and measurements vary by workload.

3.3 Understanding query results

Result	Interpretation
time	Timestamp of the recorded measurement. Results are displayed in the timestamp format returned by InfluxDB/Explorer.
host	Name of the system that supplied the measurement; the built-in source identifies the Code Creator appliance host.
usage_active	CPU activity metric from the selected CPU measurement/series.
No rows	Check the selected database and time window, then confirm the dashboard is receiving fresh data. Do not create an arbitrary new server connection as a first troubleshooting step.

3.4 Explore, visualize and manage

Explorer provides query, visualization, write and database-management functions. Use the available visualization controls on your Explorer page after a successful query. Actual controls may vary by selected Explorer view; the essential verified workflow is opening the preconfigured server, choosing telemetry, running a SQL query and seeing results.

Security Explorer is protected by the Code Creator authentication gateway on HTTPS port 8443. Do not publish the raw InfluxDB service port 8181 to make Explorer work.

4 Add Telemetry Source: optional remote monitoring

The Add Telemetry Source shortcut is for another Linux or EC2 server. It is not an initial-setup prerequisite: the appliance monitors its own EC2 host automatically through bundled Telegraf.

4.1 When to use this page

Open Add Telemetry Source when you want a separate supported Linux/EC2 host to send metrics to the appliance. Examples include an application server or a second EC2 machine that you administer.

4.2 Customer workflow

1. In Command Center choose Add Telemetry Source.
2. Read the connection requirements and supported instructions shown on that page for the server you plan to monitor. Use the source-specific instructions rather than assuming its settings match the appliance host.
3. Complete the guided source registration and collector connection steps provided by the page. Keep any generated connection credentials private.
4. Use the page's verification workflow to confirm that the remote source begins reporting. Then check its incoming measurements in Explorer.

Do not confuse the two source types Built-in host telemetry starts automatically. The System Health Check may still say no remote Linux/EC2 sources are registered until you add one. This does not by itself indicate the appliance telemetry has stopped.

The precise field names, generated remote installation command and collector network requirements are determined by the Add Telemetry Source page on your instance. They are not reproduced here because those exact fields were not part of the validated page record used for this manual.

5 System Health Check: button and result guide

5.1 Run Health Check

1. From the home page click Run Health Check. The page initially shows READY and NOT RUN for live checks.
2. Click the Run Health Check button on that page. Wait for the results to display. The button performs the implemented live database and service checks.
3. Review each section separately; an overall ATTENTION status can reflect the absence of optional remote sources rather than an appliance failure.

5.2 Interpret status labels

Label	What it means
PASS	The specific live check succeeded.
ATTENTION	Review the described condition; on a fresh launch, no registered remote sources is expected.
FAIL	The tested operation failed; review its explanation and follow the next action.
NOT RUN	The check has not been initiated yet.
STAGED	The section is visible but its restricted live backend check is not yet implemented in this release; do not treat it as PASS.

5.3 What each section checks

Page section	Validated behavior / customer meaning
Database & Data Flow	Live: database reachable, synthetic write succeeds and query returns the expected value. A PASS confirms the end-to-end tested write/query path.
Connected Data Sources	May show ATTENTION when no optional remote Linux/EC2 source is registered; built-in host telemetry is separate.
Storage	Displayed as STAGED in v1.9. Do not treat the listed writable-directory, ownership and disk-headroom items as completed live checks.
Backup & Recovery Readiness	Displayed as STAGED on this health page. Check actual backup and restore state on the dedicated Backup & Recovery page.

Page section	Validated behavior / customer meaning
Appliance Services	Live: Command Center request path, credential helper and required configuration. A PASS means those tested components responded.

If the dashboard looks healthy but the health banner says **ATTENTION** First check whether the reason is simply “No active remote Linux or EC2 data sources are currently registered.” Adding a remote source is optional.

6 AWS Configuration Settings: optional S3 authorization

Choose AWS Configuration Settings on the home page, or use Create Role from Configure Backup. This page explains how to authorize protected backups using a customer-owned private S3 bucket and an EC2 instance IAM role. It is not required to use the database or Explorer.

6.1 Step-by-step role setup shown by the page

1. Create or choose a private, general-purpose S3 bucket. Keep Block Public Access enabled, turn on bucket encryption and prefer bucket-owner-enforced object ownership / ACLs disabled. Versioning is optional.
2. In AWS IAM choose Roles > Create role; select AWS service and EC2 as the use case. Give the role a descriptive name.
3. Add a bucket-scoped permissions policy using the copyable JSON template below. Replace YOUR_BACKUP_BUCKET_NAME in BOTH Resource entries with your bucket name.
4. In the EC2 console select the InfluxDB instance, then Actions > Security > Modify IAM role. Attach the new role.
5. Return to Configure Backup and click Refresh AWS Status. Look for IAM ROLE READY, then proceed to the S3 destination validation.

6.2 EC2 role trust policy — Copy Trust Policy JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": { "Service": "ec2.amazonaws.com" },
      "Action": "sts:AssumeRole"
    }
  ]
}
```

6.3 S3 permissions — Copy Policy JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "CodeCreatorInfluxBackupBucket",
      "Effect": "Allow",
      "Action": ["s3:GetBucketLocation", "s3:ListBucket"],
      "Resource": "arn:aws:s3:::YOUR_BACKUP_BUCKET_NAME"
    },
    {
      "Sid": "CodeCreatorInfluxBackupObjects",
      "Effect": "Allow",
      "Action": ["s3:GetObject", "s3:PutObject", "s3:DeleteObject"],
      "Resource": "arn:aws:s3:::YOUR_BACKUP_BUCKET_NAME/*"
    }
  ]
}
```

Permissions Use the policy that matches the bucket your appliance will actually use. Avoid broad AmazonS3FullAccess when bucket-scoped permissions suffice. The standard policy does not include permissions for a customer-managed KMS key; add those only if your bucket requires them.

7 Backup & Recovery: every page and action

Open Backup & Recovery from Command Center. This page displays the backup destination, schedule, latest backup, guided recovery entry point and separate recovery-verification result.

7.1 Backup & Recovery overview

Area or button	What the customer should understand
BACKUP READY / destination status	Shows whether the configured S3 destination is ready for guarded backup execution.
Backup Destination	Identifies whether Amazon S3 has been configured and verified; active InfluxDB data stays on local EBS.
Backup Schedule	Shows whether an automatic schedule is saved. The scheduler evaluates the saved schedule and runs the guarded backup worker when due.
Backup Status	Shows the most recent backup ID, completion time, archive size, database pause, delivery verification and S3 object path.
Configure Backup	Opens destination, IAM status, schedule and write/read/delete testing.
Run Backup Now	Starts an immediate guarded backup to the configured S3 destination.
Start Guided Recovery	Opens the replacement-instance recovery workflow. It is not an in-place restore button for the machine that created the selected backup.
Recovery Verification	Shows whether a separate replacement appliance actually restored and validated a selected backup.

7.2 Configure Backup page and buttons

1. On Configure Backup, enter your Amazon S3 Bucket name and, if desired, an Optional S3 Prefix. A prefix groups backup objects inside your bucket; leave it blank if you do not need one.
2. In AWS Authentication, an unconfigured customer instance shows ACTION REQUIRED. Use Create Role for setup instructions; attach a role and click Refresh AWS Status until the status shows IAM ROLE READY.
3. Choose the available Backup Schedule frequency (Daily, Every 12 hours, or Weekly) and a Backup Window (UTC) suited to your workload.
4. Choose Save & Test Backup Configuration. The validation creates a temporary S3 object, reads it back, verifies its contents and deletes it. The selected destination and schedule are saved when the operation succeeds.
5. Confirm BACKUP DESTINATION VERIFIED / Validation VERIFIED. Then return to Backup & Recovery and use Run Backup Now to create your first actual backup.

Difference between tests Save & Test Backup Configuration proves temporary-object access to S3. Run Backup Now creates an appliance backup. Neither by itself proves a replacement instance can restore the database.

7.3 Reading Backup Status

Backup ID uniquely identifies a backup point. Completed is the backup time; Archive Size is the compressed archive size; Database Pause is the short reported pause during backup. Backup Verified YES refers to successful backup delivery and verification in S3. Recovery Verified is separate and requires a tested restore. The S3 object path helps identify the saved archive in your own bucket.

8 Guided Recovery: safely restore on a replacement

Guided Recovery is for a fresh compatible Code Creator appliance, not the EC2 instance that created the selected backup. Use a separate replacement EC2 instance and an IAM role authorized to read the original backup bucket. A restore changes the replacement database: do not use a target holding data you need to keep.

8.1 Page-by-page recovery workflow

1. Step 1 — Launch Replacement Appliance: launch a new compatible AMI and finish first boot using its own EC2 key pair and credentials.

2. Step 2 — Verify Replacement Instance: attach an IAM role with access to the existing backup bucket and configure/validate the same destination in Configure Backup. The page should identify the current and source instance IDs.
3. Step 3 — Select Backup: choose the exact verified backup you intend to restore. Check Backup ID, creation time, InfluxDB version, archive size and the source instance. Do not assume the newest backup is the intended one.
4. Review the preflight safety results. The page checks Replacement Instance, Appliance Family, Appliance Build, InfluxDB Version, EC2 IAM Role, S3 Backup Access, Storage Capacity and Fresh Replacement Target. Begin only when RECOVERY PREFLIGHT PASS is displayed.
5. Step 4 — Restore Appliance Data: click Begin Guided Recovery ONCE. The guarded worker performs its controlled restore and validates the resulting appliance state. Do not initiate another backup, manual restart or another restore during the operation.
6. IMPORTANT — The page does not automatically redirect to the result. Click “← Back to Backup & Recovery” at the TOP of the Guided Recovery page to inspect the recovery status. Refresh that page if the operation is still running.
7. Step 5 — Verify Recovery: find the Recovery Verification section on Backup & Recovery. Require RECOVERY VERIFIED with the expected backup ID and replacement instance ID before considering the restore complete.
8. Open Command Center > Open InfluxDB 3 Explorer on the replacement instance. Confirm the default server and telemetry database appear without manually adding a server. Query historical telemetry and verify new measurements resume.

Two different verification fields A new backup made by the replacement instance may show Recovery Verified NO even though the separate Recovery Verification section says RECOVERY VERIFIED for the older backup you restored. Always compare the relevant backup IDs.

8.2 Read-only example: check restored CPU history

In Explorer select telemetry and query a timestamp window that precedes the replacement instance's launch. Substitute dates within the history covered by your selected backup:

```
SELECT time, host, usage_active
FROM cpu
WHERE time >= '2026-09-23T16:56:00Z'
AND time < '2026-09-23T16:57:16Z'
ORDER BY time DESC
LIMIT 10;
```

The timestamps above illustrate a historical window from the Code Creator validation exercise; they will not normally contain data in a customer's own backup. Use your actual prior telemetry timestamps.

9 Administration, troubleshooting and safety

9.1 Common tasks

Task	What to do / where
Retrieve generated credentials	SSH to your EC2 instance as ubuntu; run <code>sudo cat /root/FIRST_LOGIN.txt</code> . Keep the output private.
Check services (read only)	On the EC2 instance: <code>cd /opt/codecreator/influxdb3 && sudo docker compose ps</code>
Check first boot	On the EC2 instance: <code>sudo systemctl status codecreator-influx-firstboot.service --no-pager</code>
Open the dashboard	In your browser: <code>https://YOUR_PUBLIC_IP/</code> (port 443).
Open Explorer	Click Open InfluxDB 3 Explorer from Command Center, or use <code>https://YOUR_PUBLIC_IP:8443/</code> after signing in.
Check a backup or restore	Use the dedicated Backup & Recovery page; compare the backup ID with the Recovery Verification section.

9.2 What to do when something does not look right

Symptom	First check / safe next action
Cannot reach Command Center	Check current EC2 public IP, security-group TCP 443, instance running state and first-boot/Compose status.
Browser certificate warning	Read the certificate status in FIRST_LOGIN.txt. Verify that you are using the current instance IP and that certificate issuance/renewal can complete.
Explorer says Add New Server	Do not create arbitrary connections or repeat recovery. Inspect actual recovery status; contact Code Creator support if the default server did not initialize.
No CPU query rows	Confirm telemetry database and recent time window; check dashboard RECEIVING indicator and run the supported health check.
Health Check ATTENTION on fresh launch	No optional remote Linux/EC2 sources may be registered. Built-in host telemetry is independent.
Backup role not ready	Check EC2 instance role attachment, bucket-scoped IAM policy, bucket name and region, then Refresh AWS Status.
Recovery not yet verified	Return to Backup & Recovery, refresh and inspect Recovery Verification. Do not click Begin Guided Recovery a second time.

9.3 Security and operating boundaries

- Use HTTPS for Command Center and Explorer. Limit inbound SSH (22), Command Center (443) and Explorer (8443) to trusted networks where practical.
- Keep internal InfluxDB port 8181 private. Do not paste passwords, API tokens, session cookies, IAM credentials or SSH private keys into support messages or public screenshots.
- Active database data is stored on EC2 EBS. S3 backup and recovery are optional separate operations; configure retention, storage capacity and access according to your organization's requirements.
- This release is a single-node InfluxDB Core appliance. Do not assume clustering, automatic failover or automatic migration from InfluxDB 2.x.
- EC2, EBS, S3, network-transfer and other AWS infrastructure charges are separate from any Code Creator Marketplace software charge.

For additional documentation See:

- <https://docs.influxdata.com/influxdb3/explorer/>
- <https://docs.influxdata.com/influxdb3/core/>