veeam

Mastering Veeam Backup & Replication 13

An authoritative guide to architect a resilient Veeam 13 backup platform with proven best practices





Chris Childerhose | Rick Vanover

Foreword by Nikola Pejková, Senior Community Manager at Veeam

packt

Mastering Veeam Backup & Replication

Fourth Edition

An authoritative guide to architect a resilient Veeam 13 backup platform with proven best practices

Chris Childerhose

Rick Vanover

‹packt›

Mastering Veeam Backup & Replication

Fourth Edition

Copyright © 2026 Packt Publishing
All rights reserved. No part of this book may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, without the prior written permission of the publisher, except in the case of brief quotations embedded in critical articles or reviews.
Every effort has been made in the preparation of this book to ensure the accuracy of the information presented. However, the information contained in this book is sold without warranty, either express or implied. Neither the author, nor Packt Publishing or its dealers and distributors, will be held liable for any damages caused or alleged to have been caused directly or indirectly by this book.
Packt Publishing has endeavored to provide trademark information about all of the companies and products mentioned in this book by the appropriate use of capitals. However, Packt Publishing cannot guarantee the accuracy of this information.

Portfolio Director: Vijin Boricha
Relationship Lead: Prachi Sawant
Project Manager: Gandhali Raut
Content Engineer: Swathi Ajayakumar
Technical Editor: Nithik Cheruvakodan
Indexer: Tejal Soni
Proofreader: Swathi Ajayakumar
Production Designer: Jyoti Kadam
Growth Lead: Ankita Thakur

First published: February 2021
Second edition: February 2022
Third edition: April 2023
Fourth edition: April 2026

Production reference: 1270426

Published by Packt Publishing Ltd.
Grosvenor House
11 St Paul's Square
Birmingham
B3 1RB, UK.

ISBN 978-1-80638-889-9
www.packtpub.com

I would like to dedicate this book to the Veeam 100 which inspires me daily

– Chris Childerhose

Foreword

As the Veeam Data Platform reaches its 13th version, this 4th edition of our reference book stands as the most comprehensive guide yet for a broad audience—from backup administrators who manage VDP daily to anyone seeking to master deployment, operation, and the full potential of Veeam's expanding platform. With each new release, Veeam continues to set the standard for data protection, resilience, and innovation, and V13 is no exception, introducing advanced features, expanded integrations, and enhancements in AI-driven operations that empower users to safeguard and manage their data more effectively than ever before.

This book is meticulously crafted to serve as an essential resource for all Veeam users. Whether you are navigating V13 for the first time or looking to deepen your expertise, you'll find invaluable best practices, practical real-world examples, and deep dives into the features and capabilities that matter most. The clear guidance offered here ensures that you can maximize the use of Veeam Data Platform, leveraging everything from its robust backup functionalities to the latest AI-powered advancements.

The author, Chris Childerhose, is a long-standing and highly respected member of the Veeam community. Holding both Veeam Vanguard and Veeam Legends titles, maintaining up-to-date certifications, and contributing tirelessly to the community—all while excelling in his professional career—Chris exemplifies the passion and dedication that drive Veeam's collective success. Writing a book of this scope is a significant undertaking, demanding unwavering focus and a commitment to accurately reflect the evolving experience of Veeam users. This edition highlights not just what's new, but also what truly enhances and empowers your daily work in data protection.

One of Veeam's strength lies in its vibrant community, united by a shared commitment to supporting one another, sharing the expertise and advancing the platform. This book is a testament to that spirit: it has benefited from technical reviews by Markus Hartmann, a fellow Veeam Vanguard, and Arron Stebbing, a technical evangelist with deep Veeam expertise, and includes a chapter authored by Rick Vanover, Veeam's Vice President of Product Strategy and founder of our global community.

We are proud at Veeam to count such passionate, talented individuals among our ranks. Their efforts ensure that you, the reader, have the insights and tools needed to secure, manage, and optimize your organization's most valuable asset—its data.

I wish you an enlightening and productive journey through this book and with Veeam Data Platform V13.

Nikola Pejková

Senior Community Manager, Veeam

Contributors

About the authors

Chris Childerhose is an Information Technology professional with 30+ years of experience in Network/Systems Architecture, Network & Systems Administration and Technical Support. He is a Veeam Vanguard/Veeam Legend and is a Veeam Certified Architect and Veeam Certified Engineer. He also has the following certifications: vExpert, VCAP-DCA, VCP-DCV, MCITP. He currently works for ThinkOn as the Enterprise Architect, where he designs the infrastructure for all client services offered. Chris is also an avid blogger on all things virtual, focusing on Veeam and VMware.

I would like to extend my thanks to my family for accepting that I needed to work long hours on this book during family time. I would like to thank Packt for giving me the opportunity to write this fourth edition, which was a great experience. Special thanks to Rick Vanover, and Nikola Pejkova for their contributions to this book and to the entire Veeam 100 for their support during the course of writing this book.

Rick Vanover is Vice President, Product Strategy at Veeam. Rick's passion for challenges led to his commitment to educate and communicate at all levels—engaging those new to availability technologies as well as those who are experts. As an active member of the IT community, Rick builds relationships and spreads excitement about Veeam solutions. Before becoming the "go-to" guy for Veeam questions, Rick was in system administration and IT management. He has held community designations including Microsoft MVP, VMware vExpert and Cisco Champion.

About the reviewers

Markus Hartmann is a Senior Consultant at FREICON GmbH & Co. KG and co-owner of 4IT GmbH, a consultancy specializing in data center solutions. With more than 25 years in IT and nearly 20 years of deep expertise in Veeam and VMware, he is a recognized expert in backup, virtualization, and data center architecture.

He holds leading industry certifications, including Veeam Certified Architect (VMCA), Veeam Certified Engineer (VMCE), and VMware Certified Implementation Expert (VCIX), as well as several Dell certifications.

In 2024, he was inducted into the Veeam 100 program as a Veeam Legend and has been recognized as a Veeam Vanguard since 2025. He also regularly shares his expertise on his blog: `https://markushartmann.blog`

When not working in the data center, Markus enjoys hiking the trails near his home in Zwingenberg, Germany, with his Labrador Retriever—or spending quality time with his three children.

Arron Stebbing is Founder and Director of QD Solutions Group, a consultancy focused on embedding leadership, capability, and technical expertise into organizations. With more than 20 years of experience across the Asia–Pacific region, he specializes in cloud architecture, infrastructure strategy, and technical enablement.

Arron works with enterprises, service providers, and technology vendors to design resilient cloud and data strategies, strengthen operational capability, and accelerate digital transformation. He is particularly focused on helping teams translate strategy into execution and build sustainable, high-performing environments.

Arron has extensive experience with Veeam Backup & Replication, including designing and reviewing backup and disaster recovery solutions that ensure data availability, ransomware resilience, and business continuity across hybrid and multi-cloud environments.

He is committed to simplifying complex technology challenges and enabling meaningful business outcomes.

Table of Contents

Preface **xxix**

Free benefits with your book **xxxiii**

Part 1: Installation – Veeam Software Appliance, JeOS & Veeam Backup Console **1**

Chapter 1: Installing and Upgrading Veeam **3**

Technical requirements **3**

Understanding the best practices and optimizations for Veeam installation **4**

Installing Veeam backup and replication v13 • 4

Configuring and optimizing proxy servers **20**

Setting up repository servers for success **25**

Understanding the Scale-Out backup repository **27**

VeeaMover • 37

SOBR rebalance • 38

Upgrading Veeam backup and replication to v13 **38**

Summary **44**

Further reading **45**

Chapter 2: Veeam Software Appliance (VSA) **47**

Technical requirements **47**

Introducing the Veeam software appliance **48**

Installing Veeam backup and replication 13 with OVA and ISO • 50

Security hardening and high availability **58**

Security hardened out of the box • 59

Just enough operating system (JeOS) architecture • 59

DISA STIG compliance • 59

Secure remote access • 59

Multi-factor authentication and security officer role • 61
Zero trust architecture • 62
Automated patching and updates • 62
Security features summary • 63
High availability • 63
Cluster architecture • 64
Switchover and failover operations • 65
Integration with veeam ONE • 66
Benefits of high availability • 67
Web-based management interface and enhanced RBAC..... 68
The new web-based user interface • 68
Key features and capabilities • 68
Enhanced Role-Based access control (RBAC) • 70
Built-in roles • 70
Custom role creation • 71
Examples of role configurations • 72
Security benefits of enhanced RBAC • 74
Centralized update management and auto-updates..... 74
Understanding the update architecture • 75
Update categories and policies • 76
Mandatory updates • 76
Optional updates • 76
Global update settings • 77
Component-specific configuration • 80
Configuring updates for air-gapped environments • 81
Update types summary • 82
Benefits of centralized update management • 82
Best practices for update management • 84
Summary 85
Further reading 85

Chapter 3: Just Enough Operating System (JeOS) 87

Technical requirements **87**
Veeam infrastructure appliance (JeOS) **88**
What is Veeam JeOS? • 88
JeOS vs Veeam software appliance • 88
Deployable infrastructure roles • 89
Key benefits of JeOS deployment • 89
Deploying the Veeam infrastructure appliance • 91
System requirements • 91
Step 1: Download the VIA ISO • 91
Step 2: Prepare the target system • 91
Step 3: Boot and select role • 92
Step 4: Automated installation • 93
Step 5: Initial configuration wizard • 94
Step 6: Complete configuration • 95
Managing JeOS appliances • 96
Veeam proxy and hardened repository – JeOS deployment **97**
Understanding VIA installation options • 97
Part 1: deploying a VMware backup proxy • 98
System requirements for proxy deployment • 98
Installing the infrastructure appliance for proxy role • 99
Adding the VIA server to Veeam backup and replication • 100
Assigning the VMware backup proxy role • 101
Part 2: deploying a hardened repository • 102
System requirements for a hardened repository • 102
Installing the hardened repository • 103
Adding the hardened repository to veeam backup and replication • 104
Configuring the backup repository with immutability • 105
Immutability considerations • 106
Quick reference: VIA deployment summary • 106
Out-of-the-box security hardening with JeOS ISO **107**

DISA STIG compliance by default • 107

Key DISA STIG categories implemented • 108

Zero trust architecture and privilege separation • 108

Designated security officer role • 109

Built-in immutability and ransomware protection • 109

Backup infrastructure lockdown mode • 109

Automated security updates with veeam updater • 109

Update scheduling and control • 110

Minimized attack surface • 110

SSH disabled by default • 110

Security benefits – summary • 110

Operational implications • 111

Hardened repository immutability features 111

How immutability works • 112

Block generation: optimizing immutable operations • 112

How block generation works • 112

Benefits of block generation • 113

Block generation periods by storage type • 113

Planning storage capacity with block generation • 114

XFS fast clone and block clone integration • 114

How fast clone works • 114

XFS configuration requirements • 115

GFS backups and immutability • 115

Best practices for immutability configurations • 116

Summary 116

Further reading 117

Chapter 4: Veeam Backup & Replication Console 119

History of Veeam Backup & Replication user interface 119

Deployment of the Veeam console 121

Navigating the Veeam console 123

Home section • 124

Main menu • 124

Jobs view • 129

Backups view • 130

Last 24 Hours View • 132

Inventory view • 133

Backup infrastructure view • 134

Files view • 136

Updating the Veeam Backup & Replication user interface 137

Web UI vs. Windows UI in v13 138

Advanced tips and tricks for using Veeam Console 141

Best practices for backup consoles • 141

Session statistics • 142

Text editing, FAST • 144

Secret CDP bottleneck information • 146

Backups information • 147

Hello CTRL • 147

Refresh first, avoid questions later • 148

Summary 149

Further reading 150

Part 2: Veeam AI, Best Practices for Active Directory, High Availability (HA) & Veeam CDP Enhancements 151

Chapter 5: Veeam AI – Veeam Intelligence 153

Technical requirements 153

Veeam Intelligence – A new way to help 154

Understanding Veeam Intelligence • 154

What's new in version 13 • 155

Malware threat analysis agent • 155

Deep data analysis agent • 156

Data resilience daily summary (morning coffee report) • 156

Enhanced visualization and reporting • 158

Voice interaction with speech-to-text • 158
Transparent AI processing • 158
Summary of version 13 enhancements • 158
Benefits of using Veeam Intelligence • 159
Accelerated troubleshooting and root cause analysis • 159
Proactive risk management • 159
Simplified compliance and reporting • 160
Improved cross-team communication • 160
Reduced threat dwell time • 160
Enhanced operational efficiency • 161
Practical use cases • 161
Accessing Veeam Intelligence 161
Prerequisites • 162
Accessing Veeam Intelligence in web UI • 162
Launching Veeam Intelligence • 162
Using the chat interface • 163
Configuring Veeam Intelligence settings via web UI • 165
Accessing Veeam Intelligence in Windows console • 166
Launching Veeam Intelligence • 166
Managing the chat window • 166
Using the chat interface • 167
Configuring Veeam Intelligence settings via console • 168
Best practices • 168
Security considerations • 169
Troubleshooting • 169
Veeam Intelligent operating modes 170
Understanding the operating modes • 170
Disabled mode • 170
Basic mode • 171
Advanced mode • 171
Malware deep analysis agent • 172
Code generation and analysis • 172

Veeam Intelligence: AI-powered assistance for backup administrators 173
Simplifying daily operations • 173
Capacity planning and forecasting • 174
Compliance and reporting • 174
Knowledge assistance and best practices • 175
Practical considerations • 175
Summary 176
Further reading 176

Chapter 6: Best Practices for Active Directory Domains with Veeam 177

Active Directory integration with Veeam Backup & Replication 177
Understanding the authentication model • 178
Prerequisites for Active Directory integration • 178
Configuring role-based access control with AD groups • 179
Service account configuration and delegation • 180
Security considerations and best practices • 180
Troubleshooting common integration issues • 181
Integration with Veeam software appliance • 182
Best practices for Active Directory with Veeam Backup & Replication v13 183
Veeam infrastructure and Active Directory: Domain placement strategy • 184
Understanding Active Directory forests and domains • 184
Deployment options • 185
Windows workgroup deployment • 186
Workgroup advantages • 186
Workgroup disadvantages • 186
Management domain deployment • 187
Management domain advantages • 187
Management domain disadvantages • 188
Deployment model comparison • 188
Forest trust configuration • 189
Authentication and identity management in v13 • 191
SAML 2.0 single sign-on • 191

Enhanced role-based access control • 191
Kerberos vs. NTLM authentication • 192
Leveraging Active Directory for agent-based protection groups • 193
Security recommendations summary • 193
Securing Active Directory for Veeam Backup & Replication v13 194
Deployment models: Workgroup, management domain, or production domain • 195
Option 1: Separate management forest (recommended) • 195
Option 2: Workgroup deployment • 197
Option 3: Production domain (hardened) • 198
Securing service accounts • 199
Group managed service accounts (gMSA) • 199
Standard service account hardening • 200
Multi-factor authentication (MFA) • 201
Enabling MFA in the Veeam console • 201
MFA for the Veeam software appliance host management UI • 202
Network segmentation and DNS considerations • 203
Summary 204
Further reading 204

Chapter 7: Veeam High Availability (HA) 207

Technical requirements 207
High Availability: Protecting your environment for disaster recovery 208
Why HA matters for disaster recovery • 208
Overview of HA cluster architecture • 209
Core components • 209
Data synchronization and replication • 210
Split-brain prevention • 211
Disaster recovery design patterns with HA • 211
Pattern 1: Same-site HA for hardware and OS failures • 211
Pattern 2: Cross-site HA for site-level disasters • 211
Pattern 3: Combined HA with Veeam ONE automated failover • 212
Alternatives when HA cannot be deployed • 212

High Availability (HA) cluster 213
Understanding High Availability • 213
How the HA cluster works • 215
Cluster assembly • 215
Data synchronization • 215
WAL file storage and resilience • 216
Infrastructure requirements • 216
Licensing • 216
Server requirements • 216
Network and DNS configuration • 217
HA cluster internal components • 217
Configuring the HA cluster • 218
Phase 1: Enabling HA on both nodes • 218
Phase 2: Assembling the cluster via the VBR console • 221
Switchover and failover operations • 224
Switchover (planned) • 224
Failover (unplanned) • 226
Switchover vs. failover • 227
Post-failover recovery • 227
Managing the HA cluster • 228
Monitoring cluster status • 228
Rescanning the cluster • 228
Disassembling the cluster • 228
Automatic updates in an HA cluster • 228
Benefits of using the HA cluster • 229
Considerations and limitations • 230
High Availability failover 231
Understanding HA failover architecture • 231
Failover event sequence • 232
Automatic vs. manual failover • 233
Automatic failover • 233
Manual (planned) failover • 234

Monitoring failover events • 235

Veeam ONE alarms and actions • 235

System logs and Patroni logs • 235

Post-failover operations • 236

Validating active node • 236

Recovering a failed node • 236

Failover testing and validation • 237

Recommended testing cadence • 238

Common failover issues and troubleshooting • 238

Best practices summary • 240

Summary 240

Further reading 241

Chapter 8: Veeam CDP – Universal Continuous Data Protection 243

Technical requirements 243

What's new with Universal CDP in v13.0.1 244

The vSphere limitations • 244

What is Universal CDP? • 244

CDP coordinator service • 245

Volume filter driver: I/O interception without snapshots • 245

Replication targets in v13.0.1 • 246

VMware vSphere • 246

VMware cloud director • 246

Integration with Veeam cloud connect for DRaaS • 247

Deployment overview • 247

Licensing • 248

Current limitations and known constraints • 248

Network resilience with CDP 250

Configuring continuous data protection • 250

Overview of CDP architecture and components • 250

Core components • 251

Data flow sequence • 252

Prerequisites and pre-configuration requirements • 253
I/O filter deployment • 253
CDP proxy sizing considerations • 258
Creating a CDP policy • 262
Step 1: Launching the CDP policy wizard • 262
Step 2: Selecting source VMs • 263
Step 3: Configuring the target • 264
Step 4: Assigning CDP proxies • 265
Step 5: Defining RPO settings • 266
Step 6: Configuring network mapping • 268
Step 7: Configuring guest processing • 268
Step 8: Summary and application • 269
Improving network resilience for CDP • 271
Understanding CDP network failure modes • 272
Network traffic encryption and compression • 273
Configuring network throttling rules • 273
WAN accelerator integration • 274
Source proxy cache configuration • 275
Redundant network paths and proxy failover • 275
Monitoring CDP health and RPO compliance • 276
CDP network resilience best practices • 276
Enhanced I/O journal scalability and additional CDP features 278
I/O journal scalability enhancements • 278
Journal architecture and storage mechanics • 278
Scalability benchmarks • 279
CDP I/O filter cross-compatibility • 280
Universal CDP: Agent-based continuous data protection • 281
Agent-based architecture • 281
Replication targets • 282
Universal CDP data flow • 283
Failover and recovery enhancements • 283
I/O throughput visualization during failover • 283

Planned CDP failover • 284

File-level restore from CDP replicas • 284

CDP with Veeam cloud connect for DRaaS • 284

Design considerations • 285

Summary 286

Further reading 286

Part 3: Enhancements, Backup Infrastructure, Veeam ONE v13 & Recovery Options 287

Chapter 9: Other Enhancements – Linux Only Infrastructure and Automatic Upgrades 289

The new 64-bit architecture 289

64-bit-only: The end of 32-bit support • 290

Minimum supported operating systems • 290

gRPC: Replacing legacy communication protocols • 291

Impact on network architecture and firewall rules • 292

Authentication modernization: NTLM deprecation and certificate-based authentication • 293

The Veeam deployment kit • 293

BLAKE3 hashing: Faster, more secure data integrity • 294

Instant recovery engine and data mover improvements • 295

PostgreSQL 17.6: A unified configuration database • 295

A cross-platform architecture for the future • 296

Linux-only backup infrastructure 296

New Linux infrastructure roles in v13 • 298

Mount server • 298

Gateway server • 298

Guest interaction proxy • 299

Additional roles • 299

Veeam infrastructure appliance • 299

Certificate-based authentication and reduced port requirements • 301

Designing a Linux-only backup infrastructure 302

Core architecture components • 302
Role placement guidelines • 302
Windows console considerations • 303
Migration considerations • 304
Security benefits of a Linux-only infrastructure • 305
Centralized updates with Veeam updater 306
Architecture and components • 306
Configuring automatic updates • 307
Enabling the Veeam updater • 307
Monitoring and troubleshooting updates • 310
Update history and status • 310
Common troubleshooting scenarios • 311
Custom role-based access control (RBAC) 312
Understanding the RBAC architecture • 312
Permissions, roles, and assignments • 312
Built-in roles vs. custom roles • 313
Creating custom roles • 314
Permission categories • 319
Assigning roles and defining scope • 320
Role assignment process • 321
Scope inheritance and precedence • 321
Real-world RBAC design patterns • 322
Pattern 1: Tiered administration • 322
Pattern 2: Geographic segmentation • 322
Pattern 3: Separation of duties • 322
Auditing and compliance • 323
PowerShell integration for RBAC management • 323
Summary 324
Further reading 324

Chapter 10: Backup Infrastructure – PostgreSQL 17.6, Improved Job Manager Scalability, and Other Changes 325

PostgreSQL 17.6 support 325

PostgreSQL 17: Key improvements relevant to VBR • 326

Supported deployment configurations • 327

Fresh installation with PostgreSQL 17.6 • 327

Upgrading from previous PostgreSQL versions • 329

Configuration database maintenance • 329

Session history table growth • 329

Index bloat • 330

Connection pooling • 330

Known limitations and considerations • 330

Improved job manager scalability 331

Job manager architecture • 331

Scheduling engine improvements • 332

Parallel schedule evaluation • 332

Schedule condition caching • 332

Improved after-job chaining • 332

Dispatch layer improvements • 333

Resource utilization tracking decoupled from scheduling • 333

Proxy selection algorithm optimization • 333

Repository queue management • 333

Session state management • 334

Reduced memory footprint per active session • 334

Session history write optimization • 334

Large-scale deployment guidance • 335

VBR server sizing • 335

Backup window design • 335

Monitoring scalability health • 336

Known limitations and considerations • 336

Object storage and immutability improvements 337

Object storage enhancements • 337
Capacity tier offload performance • 337
Multipart upload reliability • 338
Archive tier improvements • 338
Selective archive tiering by retention policy • 339
Archive retrieval cost estimation • 339
Direct backup to object storage • 339
Synthetic full backup support • 339
Immutable backup chain integrity verification • 340
Immutability enhancements • 340
S3 object lock improvements • 340
Lock period alignment with retention • 340
Immutability status in the VBR console • 341
Azure Blob immutability improvements • 341
Container-level policy validation • 341
Immutability policy renewal • 341
Hardened Linux repository immutability • 342
Immutability period per backup job • 342
Immutability verification reports • 342
Scale-out backup repository operational changes • 342
Known limitations and considerations • 343
Summary 344
Further reading 344

Chapter 11: Veeam ONE v13 345

Technical requirements 345
What's new with Veeam ONE v13 346
Threat center – The new default dashboard • 346
Data platform scorecard • 347
Malware detections widget • 347
RPO anomalies widget • 347
SLA compliance overview • 348

Redesigned dashboard catalog • 348
New and enhanced alarms • 348
Application backup policy state alarm (v13.0.1) • 349
VBR high availability health alarms (v13.0.1) • 349
Coveware recon integration alarms (v13.0.1) • 349
VBR authentication anomaly alarms (v13.0.1 Patch 1) • 349
New reporting engine • 349
Key capabilities of the new engine • 350
Export format changes • 350
Redesigned report catalog • 351
Infrastructure monitoring updates • 352
Enterprise application plug-in monitoring (v13.0.1) • 353
Cloud workload monitoring expansion (v13.0.1) • 353
Veeam Backup & Replication monitoring improvements • 354
Job calendar view • 354
Credentials manager • 355
Data collection configuration in web client • 355
Veeam intelligent diagnostics updates • 356
Veeam intelligence integration • 356
Malware threat analysis agent (v13.0.1) • 357
Deep-data analysis agent (v13.0.1) • 357
Data resilience summary (v13.0.1) • 357
Veeam ONE v13 - Architecture changes 360
Veeam analytics service • 361
Veeam analytics service architecture • 361
Deployment and push installation • 362
Dual-database architecture • 363
Monitoring database (SQL server) • 364
Reporting database (PostgreSQL) • 364
Implications for administrators • 365
System requirements • 365
Operating system • 365

Hardware (All-in-one deployment baseline) • 366
Software prerequisites • 366
New monitored platform support added in v13 / v13.0.1 • 367
Network port reference • 367
Reporting engine architecture • 370
Export format changes • 370
Data flow • 371
Web client and desktop client • 371
Security architecture changes • 372
Protocol modernization: gRPC replaces WMI/RPC • 372
NTLM deprecation • 373
SAML-based single sign-on • 373
RBAC and role changes • 373
REST API authentication • 374
Integration architecture • 374
Veeam intelligence • 374
Coveware recon scanner integration • 375
Intelligent diagnostics signature updates • 376
Upgrade path from v12 to v13 • 377
Upgrade wizard overview • 377
New dashboard, Veeam intelligence, and reporting changes 381
The redesigned web client experience • 381
Dashboard catalog redesign • 381
Appearance and theming • 382
Global search • 383
Credentials manager and data collection • 383
Navigation structure • 383
Interactive dashboards and widgets • 384
Predefined dashboard highlights • 385
Heatmaps • 385
Custom dashboard creation • 386
Veeam intelligence: AI-powered observability • 386

Architecture and integration • 386
Operating modes • 387
Malware threat analysis agent • 387
Deep data analysis agent • 387
Data resilience summary • 388
The new reporting engine • 388
PostgreSQL reporting database • 388
HTML5-based interactive reports • 389
Redesigned report catalog • 389
Reporting database configuration • 390
Threat center integration and security reporting • 391
Coveware recon integration • 391
ServiceNow integration enhancements • 391
Monitoring the VSA • 391
Integration with the VBR console • 392
Key additions in version 13.0.1 • 393
Summary 394
Further reading 394

Chapter 12: Recovery Options 395

Instant recovery to Microsoft Azure 395
What is instant recovery to Azure? • 396
What's new in v13: From on-premises only to Azure-native recovery • 396
Architecture and how it works • 397
Helper appliance templates • 397
The recovery process flow • 398
Supported backup sources and repository types • 398
Prerequisites and considerations • 399
Azure account and subscription requirements • 400
Operating system and workload constraints • 400
Network considerations • 401
Instant recovery to Azure wizard • 401

Finalizing instant recovery: Migrate or discard • 403
Mass recovery considerations • 404
Universal CDP: Continuous data protection for near-zero recovery 405
From vSphere CDP to Universal CDP • 406
Universal CDP architecture and components • 407
I/O journal engine • 407
CDP proxy • 407
CDP target and replica datastores • 408
Network resilience and I/O journal continuity • 408
Recovery operations with Universal CDP • 409
Point-in-time failover • 409
Planned failover • 410
Failback and reverse protection • 410
Balancing RPO against infrastructure overhead • 411
Universal CDP in the context of the broader recovery strategy • 411
Monitoring and troubleshooting CDP health • 413
Key metrics to monitor • 413
Common troubleshooting scenarios • 413
Veeam's other recovery options 414
File-level recovery (FLR) • 415
How FLR works • 416
Windows vs. Linux FLR • 418
Volume-level restore • 419
Virtual disk restore (instant disk recovery) • 421
Restoring a virtual disk • 422
Application item restore – Veeam explorers • 425
Launching a Veeam explorer • 426
Export backup • 426
Secure restore • 427
Securing restore integration with recovery • 427
Instant cloud recovery to Microsoft Azure (v13) • 428
Bare metal recovery (BMR) • 429

BMR process overview • 429
Quick reference for recovery options • 430
Choosing the right recovery option • 431
Summary 431
Further reading 432

Chapter 13: Unlock Your Exclusive Benefits 433

Unlock this Book's Free Benefits in 3 Easy Steps..... 434

Other Books You May Enjoy 438

Index 441

Preface

Veeam Backup & Replication v13 is a core component of the Veeam Data Platform and represents a significant shift toward security-first, Zero Trust data protection. For backup administrators, mastering v13 is critical to protecting virtual, cloud, and hybrid workloads against outages, accidental data loss, and modern ransomware threats.

This book provides a practical, administrator-focused guide to designing, deploying, and operating data protection environments using Veeam Data Platform v13. It emphasizes real-world backup operations, infrastructure hardening, and recovery execution, enabling backup admins to move beyond basic job configuration toward resilient, audit-ready designs.

You will gain hands-on experience with both foundational and advanced Veeam capabilities, including continuous data protection (CDP) for low-RPO workloads, Veeam AI using Veeam Intelligence, instant recovery for rapid workload restoration, monitoring, and analytics with Veeam ONE, High Availability for disaster recovery, and a look into the new Veeam console with all its new enhancements. Throughout the book, cyber-resilience concepts such as immutability, separation of duties, least privilege, and secure recovery validation are reinforced to align with Veeam v13's Zero Trust architecture.

Who this book is for

Users will be VMware administrators or backup administrators with existing knowledge of Veeam, and the topics covered in this book. They will know virtualization and backup concepts to understand what gets discussed in each chapter. Veeam is one of the leading modern data protection solutions, and learning this technology will help users protect their environments. Most users will want to implement many of the topics discussed and investigate Veeam ONE for monitoring/reporting their infrastructure and High Availability for DR scenarios.

What this book covers

Chapter 1, Installing and Upgrading Veeam, explains how to upgrade existing environments to Veeam Backup & Replication v13 and perform clean deployments using the new appliance-based installation model.

Chapter 2, Veeam Software Appliance (VSA), covers the v13 appliance architecture, including hardened design, the new web UI, high availability, enhanced role-based access control, and security-first operations.

Technet24

Chapter 3, Just Enough Operating System (JeOS), examines the JeOS foundation used in Veeam v13, focusing on default hardening, reduced attack surface, and supported infrastructure roles.

Chapter 4, Veeam Backup & Replication Console, reviews the refreshed console experience, including UI updates, dark mode, credential management improvements, and operational enhancements.

Chapter 5, Veeam AI - Veeam Intelligence, introduces AI-driven capabilities in Veeam Data Platform v13 that provide operational insights, anomaly detection, and guidance for backup administrators.

Chapter 6, Best Practices for Active Directory Domains with Veeam, outlines secure design and operational best practices for integrating Veeam with Active Directory in Zero Trust environments.

Chapter 7, Veeam High Availability (HA), explains high-availability capabilities in v13 and how they improve resilience, simplify disaster recovery, and protect critical backup components.

Chapter 8, Veeam CDP – Universal Continuous Data Protection, explores Universal CDP, improved network resilience, and enhanced I/O journal scalability for low-RPO workloads.

Chapter 9, Other Enhancements – Linux Only Infrastructure and Automatic Upgrades, covers key architectural updates in v13, including the move to 64-bit architecture, Linux-only components, and automatic updates.

Chapter 10, Backup Infrastructure – PostgreSQL 17.6, Improved Job Manager Scalability, and Other Changes, examines backend improvements such as PostgreSQL 17.6 and increased job manager scalability for large environments.

Chapter 11, Veeam ONE v13, focuses on new monitoring, reporting, and analytics capabilities for proactive operations and security visibility.

Chapter 12, Recovery Options, concludes with modern recovery methods in v13, including instant recovery to Microsoft Azure and CDP-based recovery scenarios.

To get the most out of this book

You should have at least six months of hands-on knowledge with Windows/Linux servers and virtualization with VMware, and it would be best if you were comfortable setting up servers and configuring them with storage. You should also have some backup knowledge and have already used Veeam even for basic tasks, since many topics in the book look at the more advanced features of Veeam Backup & Replication.

Software/hardware covered in the book	Operating system requirements
Windows 2019/2022/2025	Windows
Veeam Data Platform & Veeam ONE	Windows
Veeam Software Appliance & JeOS	Linux

Windows server set up to install Veeam Backup & Replication and Veeam ONE. Veeam Data Platform and Veeam Software Appliance OVA/ISO files and trial licenses downloaded from `http://www.veeam.com` website

Download the color images

We also provide a PDF file that has color images of the screenshots/diagrams used in this book. You can download it here:`https://packt.link/gbp/9781806388899`.

Conventions used

There are a number of text conventions used throughout this book.

`CodeInText`: Indicates code words in text, database table names, folder names, filenames, file extensions, pathnames, dummy URLs, user input, and Twitter handles. For example: "Run the `setup.exe` file on the mounted ISO drive:"

Bold: Indicates a new term, an important word, or words that you see on the screen. For instance, words in menus or dialog boxes appear in the text like this. For example: "When you first launch the **Veeam Backup & Replication** console, you will be taken directly to the **Home** tab, and **Jobs** will be in focus."

Note

Warnings or important notes appear like this.

Tip

Tips and tricks appear like this.

Technet24

Get in touch

Feedback from our readers is always welcome.

General feedback: If you have questions about any aspect of this book or have any general feedback, please email us at `customercare@packt.com` and mention the book's title in the subject of your message.

Errata: Although we have taken every care to ensure the accuracy of our content, mistakes do happen. If you have found a mistake in this book, we would be grateful if you reported this to us. Please visit `http://www.packt.com/submit-errata`, click **Submit Errata**, and fill in the form.

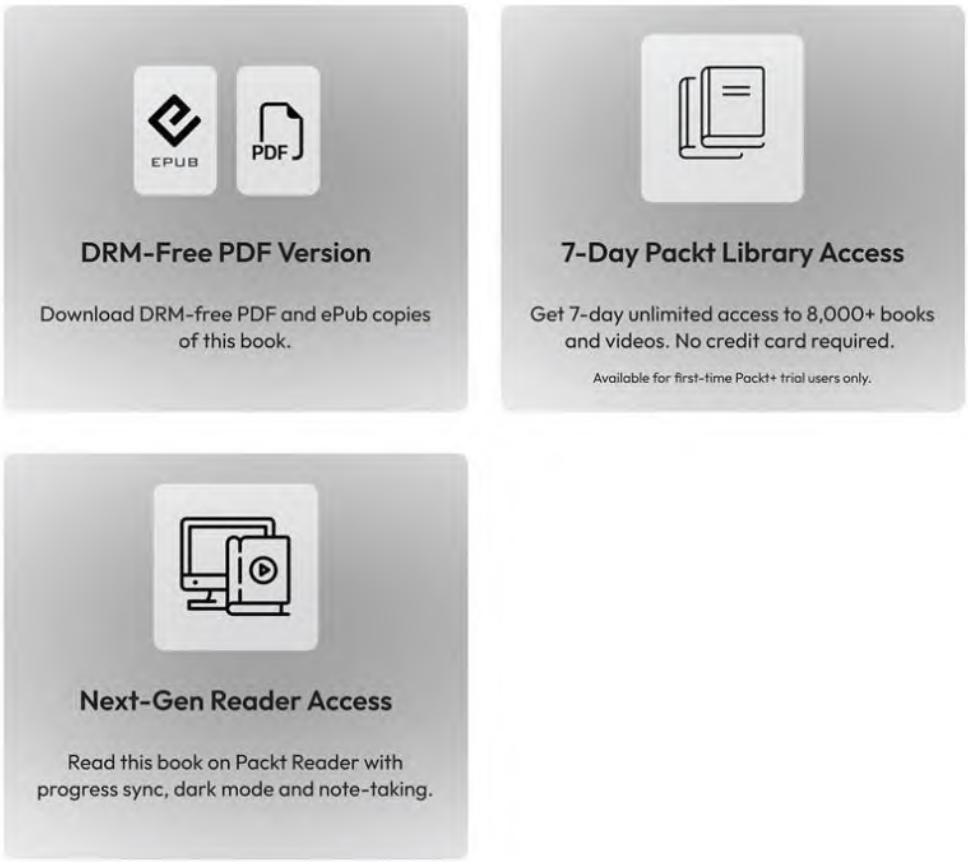
Piracy: If you come across any illegal copies of our works in any form on the internet, we would be grateful if you would provide us with the location address or website name. Please contact us at `copyright@packt.com` with a link to the material.

If you are interested in becoming an author: If there is a topic that you have expertise in and you are interested in either writing or contributing to a book, please visit .

Free benefits with your book

This book comes with free benefits to support your learning. Activate them now for instant access (see the "*How to Unlock*" section for instructions).

Here's a quick overview of what you can instantly unlock with your purchase:



Technet24

How to Unlock

Scan the QR code (or go to `packtpub.com/unlock`). Search for this book by name, confirm the edition, and then follow the steps on the page.





Note: Keep your invoice handy. Purchases made directly from Packt don't require one

Share your thoughts

Once you've read *Mastering Veeam Backup Replication*, we'd love to hear your thoughts! Scan the QR code below to go straight to the Amazon review page for this book and share your feedback.



https://packt.link/r/1806388898

Your review is important to us and the tech community and will help us make sure we're delivering excellent quality content.

Part 1

Installation – Veeam Software Appliance, JeOS & Veeam Backup Console

The objective of this section is to teach the reader best practices and optimizations for installing and upgrading Veeam. This section will also include the new Veeam Software Appliance (VSA) and Just Enough Operating System (JeOS). We will also discuss the Veeam Backup Console and what has changed in Veeam Backup & Replication v13. Readers will be able to apply the best practices and optimizations to their installation of Veeam.

This part of the book includes the following chapters:

- *Chapter 1, Installing and Upgrading Veeam*
- *Chapter 2, Veeam Software Appliance (VSA)*
- *Chapter 3, Just Enough Operating System (JeOS)*
- *Chapter 4, Veeam Backup & Replication Console*

Technet24

1

Installing and Upgrading Veeam

Veeam Backup & Replication v13 is part of the **Veeam Data Platform**, the newest release for Veeam which allows you to back up all your workloads, including **cloud**, **virtual**, **physical**, and **applications**. As we have seen from previous versions, it is simple yet flexible to meet your most challenging environment needs. This chapter will discuss how to install and upgrade the software, what components make up Veeam Backup & Replication v13, and some best practices and optimizations. There will be practical examples of optimizing specific elements that make up the **Veeam** environment throughout this chapter. We will also touch on some websites, such as the *Best Practices Guide for Veeam*, to give you resources to help set up Veeam in your environment. As they say about Veeam – "*It Just Works.*"

In this chapter, we're going to cover the following main topics:

- Understanding the best practices and optimizations for Veeam installation
- Configuring and optimizing proxy servers
- Setting up repository servers for success
- Understanding the scale-out backup repository
- Upgrading Veeam Backup and replication to v13

Technical requirements

To ensure a successful installation, you will require the following:

- You must have deployed a **Windows 2019/2022/2025** server with the necessary disk space to install the application (**2012/2008 R2 SP1** is no longer supported.) **Windows 10** and other modern **Windows** desktop operating systems are also supported. Please see the following website for the Operating System requirements: `https://helpcenter.veeam.com/docs/vbr/userguide/system_requirements.html?ver=13`

Technet24

- You must have downloaded the latest **ISO file** from `www.veeam.com`, for which you must register on the site and allow you to obtain a trial license. At the time of this writing, version `13.0.1.1071` is the latest release.
- *Veeam Best Practices* website: `https://bp.veeam.com/vbr/`.
- *Veeam Documentation* website: `https://helpcenter.veeam.com/docs/vbr/userguide/overview.html?ver=13`.

> **Note**
>
> **Your purchase includes a free PDF copy + exclusive extras**
>
> Your purchase includes a DRM-free PDF copy of this book, 7-day trial to the Packt+ library (no credit card required), and additional exclusive extras. See the *Free benefits with your book* section in the *Preface* to unlock them instantly and maximize your learning.

Understanding the best practices and optimizations for Veeam installation

Installing Veeam Backup & Replication v13 is a straightforward yet simple process. Setting up Veeam, if not done correctly, can lead to components not functioning properly and poor performance, among other issues. It will protect your data and environment with minimal configuration when set up correctly. This section will guide you through the installation process and cover best practices and optimizations for your environment.

Installing Veeam backup and replication v13

Before installing **Veeam Backup & Replication v13**, ensure that a server is deployed as outlined in the *Technical requirements* section, with sufficient disk space for the installation.

> **Note**
>
> Veeam will configure the default backup repository on the drive with the most available disk space, whether it is the OS drive, Application drive or Catalog drive.

As a best practice for installing Veeam, the disk layout should be similar to the following so that you can separate components to ensure the best possible performance:

- **OS drive**: This is where your operating system resides and should be used only for this purpose.

- **Application drive**: This will be your application installation drive for Veeam and all its components.
- **Catalog drive**: Veeam uses a Catalog that can generate around 10 GB of data per 100 VMs backed up with file indexes. If this is a significant storage requirement for your deployment, it may be advisable to allocate the *Catalog folder* to a separate drive.

Once your server is ready and you have downloaded the **ISO file** and mounted it, follow these steps to install Veeam:

1. Run the `setup.exe` file on the mounted ISO drive.

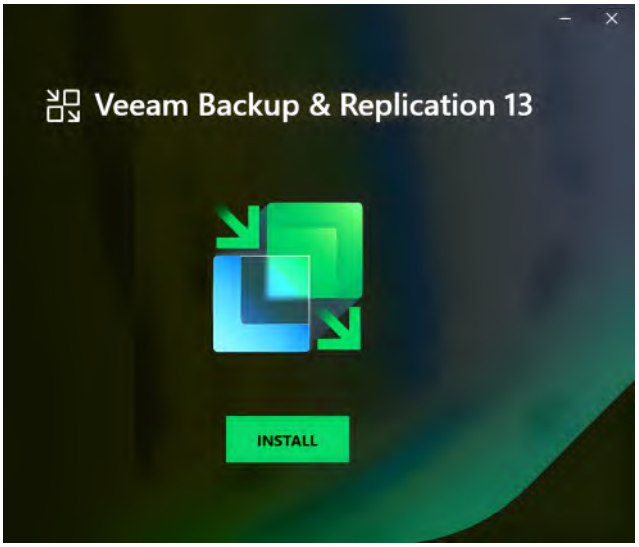


Figure 1.1: Main installation screen

2. Click the **Install** button on this screen as shown in *Figure 1.1*. You will then be presented with the component installation screen.
3. Click on the **Install Veeam Backup & Replication** as shown in *Figure 1.2*.

Technet24

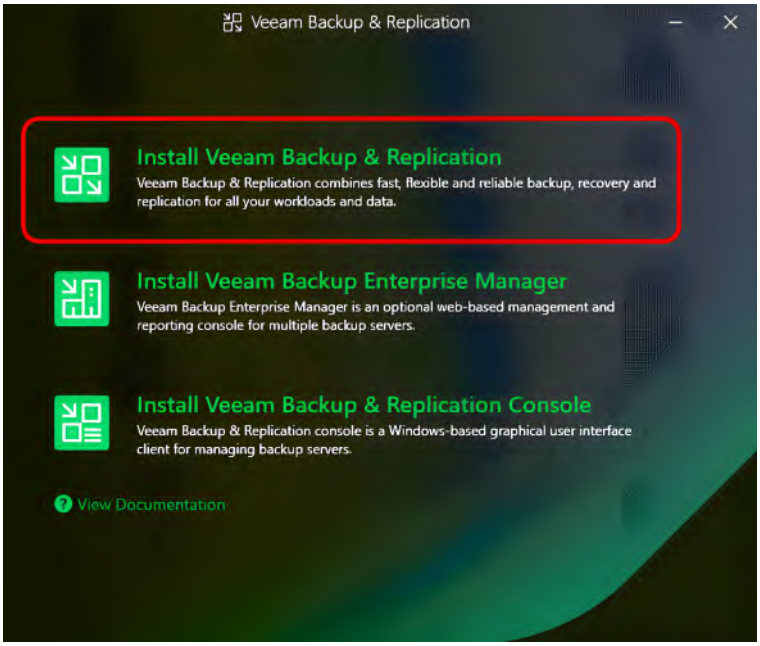


Figure 1.2: Veeam Backup and Replication installation option

4. You will see the **License Agreement** window at this point, so you must click **I Accept** to continue.

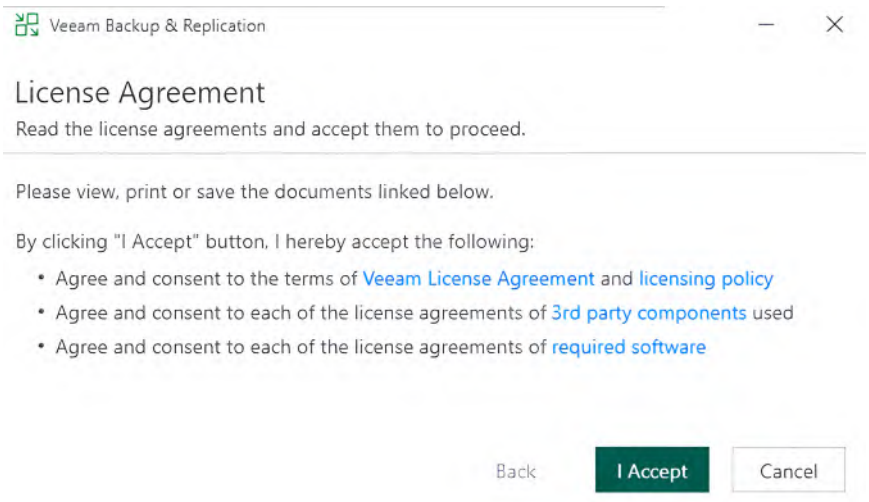


Figure 1.3: License Agreement

5. You will now need to provide a valid license file, whether purchased or a trial. If you do not have this at this stage of the installation, you can click **Next** to continue, and Veeam will operate in the *Community (Free) Edition*. When you obtain the license file, you can install it within the application under the menu and **License**.

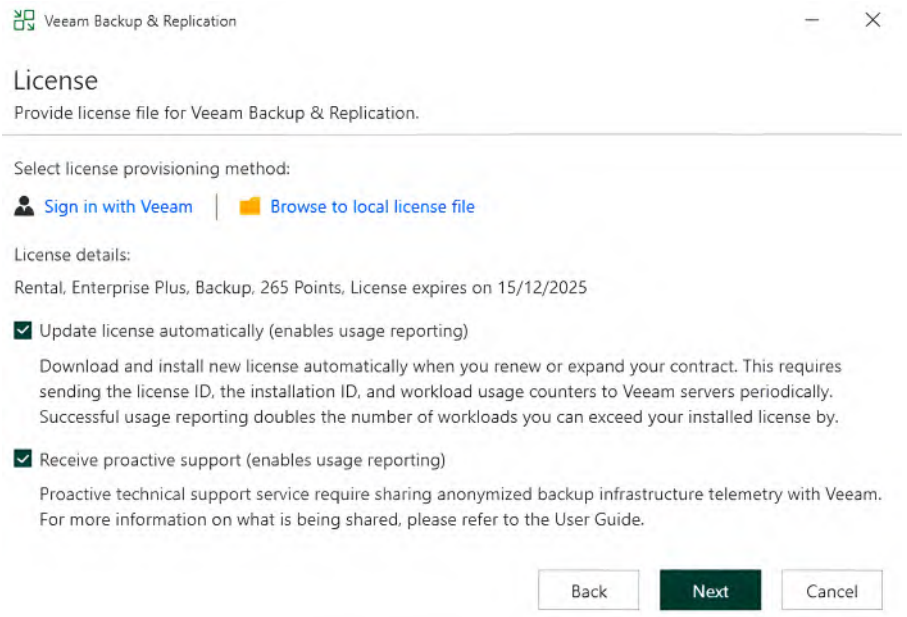


Figure 1.4: License dialogue window

Note

With Veeam Backup and Replication 13, a new checkbox is available for receiving proactive support, which you can enable if desired.

6. After clicking **Next**, the installer will do a system check for any required prerequisites. Should something be missing, you will be prompted and have the option to install the missing components.

Technet24

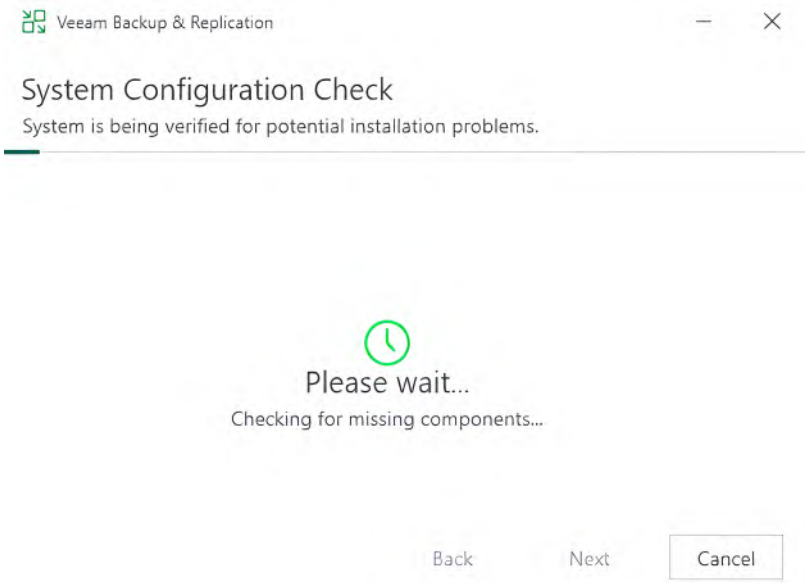


Figure 1.5: System Configuration Check – missing components

Note

An example of missing components could be a required version of Microsoft .NET or Microsoft Visual C++ Redistributable.

7. Once all the components get installed, you move to the following screen shown in *Figure 1.6*. Unlike in previous Veeam Backup & Replication versions, the screen does not allow you to input a user account to run the services. Instead, with Version 13 of Veeam, you need to select the **Customize Settings** link.

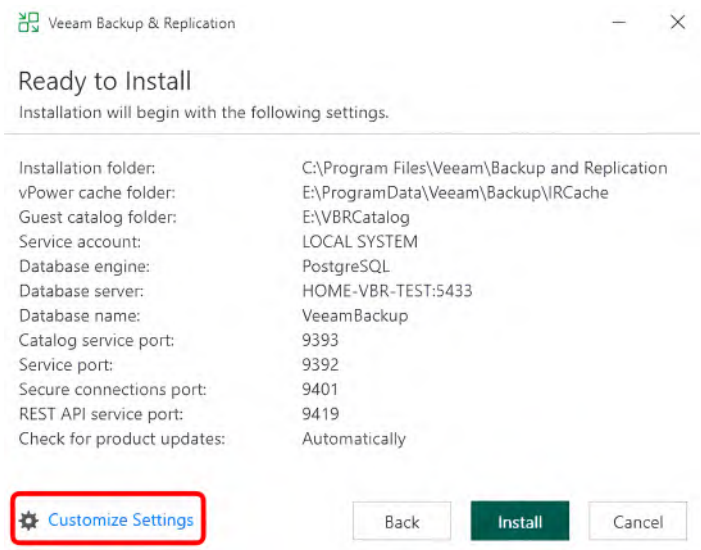


Figure 1.6: Option to specify different settings

Note

You will also notice that under the Default Configuration (Database engine), PostgreSQL Server is now the default, which we will examine once the installation is complete. If you want to use Microsoft SQL Server, you will need to have that set up ahead of time with a minimum version of Microsoft SQL Server 2016.

You will now be able to enter a user account for the Veeam services, better known as a **Service Account**. There are some recommended settings for this service account:

- You must have **Local Administrator** rights on the Veeam server.
- If you use a separate SQL Server, you will require permission to create the database. The same applies to the PostgreSQL option as well.
- You will need full NTFS permissions to the folder containing the Catalog.

Technet24

For more details about these permissions, please visit `https://helpcenter.veeam.com/docs/vbr/userguide/required_permissions.html?ver=13`.

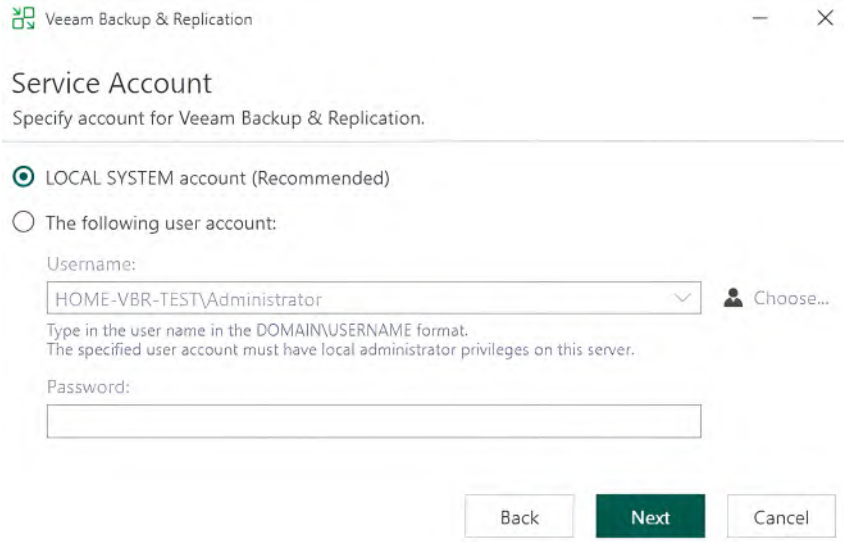


Figure 1.7: Server user account

I am using the **LOCAL SYSTEM** on my lab server for this setup. In contrast, in a production scenario, you would already have a service account set up in Active Directory to enter at this stage under **The following user account:**.

8. Now, select your database option, including the required Windows or SQL credentials. Click on **Next** to continue.

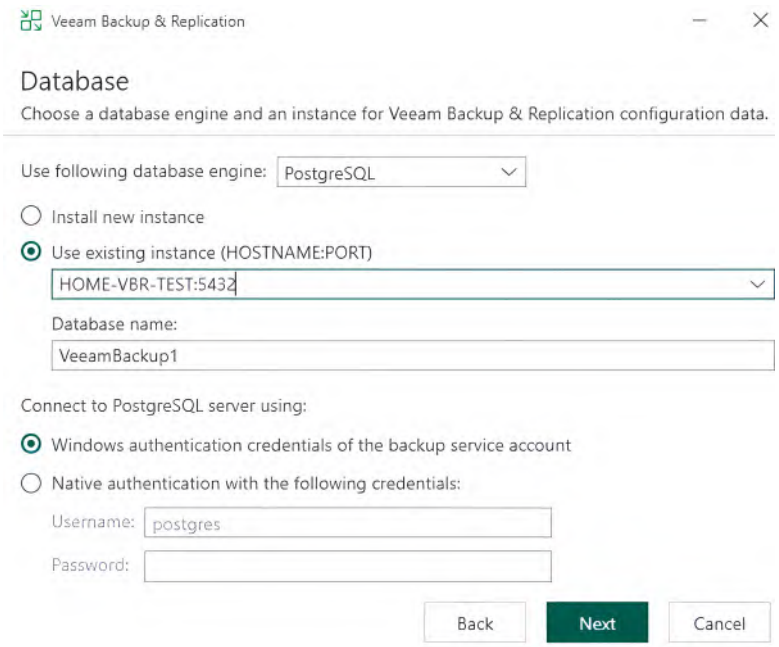


Figure 1.8: Database server selection and credentials used

9. As you can see in *Figure 1.9*, the preceding screen allows you to change the installation locations for the installation components. Click **Next** once you have completed this part of the installation.

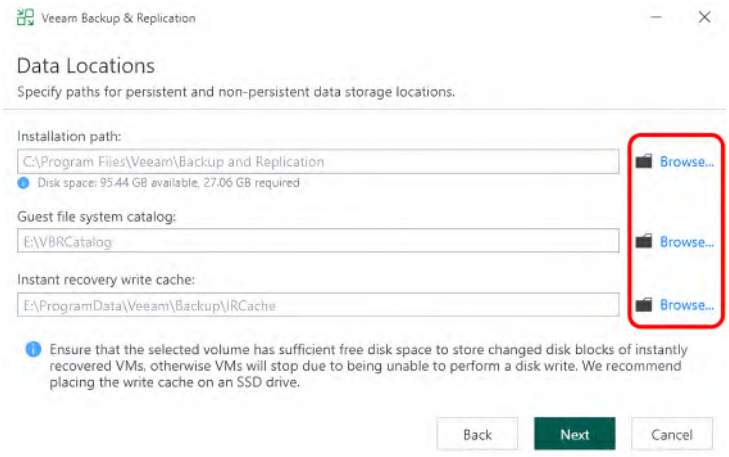


Figure 1.9: Installation paths for application

Technet24

10. You are then prompted for the **Port Configuration** to be used. Leave the defaults unchanged, then click **Next**.

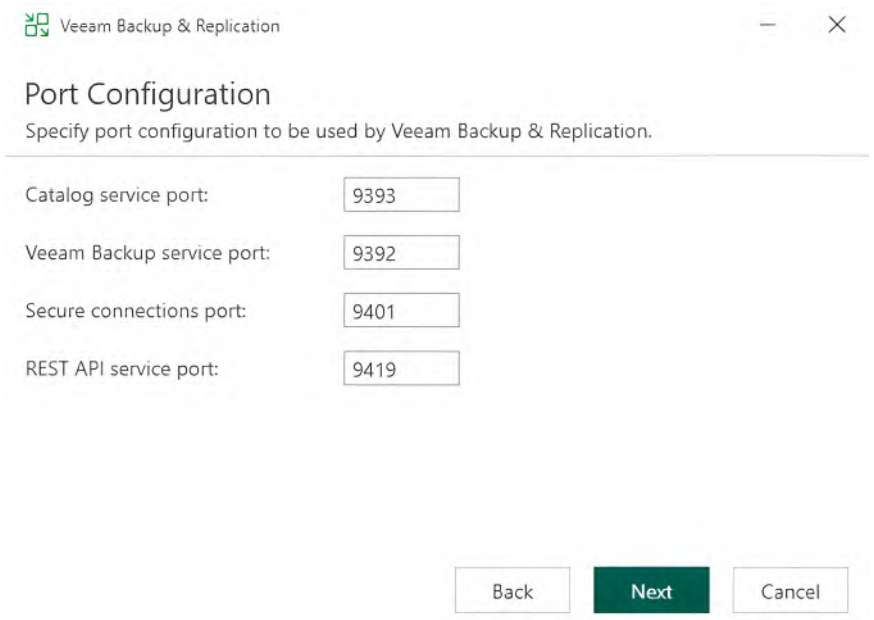


Figure 1.10: Port Configuration selection

Note

With the release of Veeam Backup and Replication v13, the default database type has been changed to *PostgreSQL*. This choice allows users to transition from Microsoft SQL, which requires licensing (SQL Express does not, but has a 10GB limit on the database size), whereas PostgreSQL does not, as it is open source. The PostgreSQL database can be installed directly on the Veeam Backup Server or on any other database server, such as those running Windows or Linux. For more details on PostgreSQL, please see the *Further reading* section at the end of the chapter.

The installer is now ready to complete the setup by installing the local PostgreSQL Server instance, followed by the application. Veeam will also set the user account you selected to start all the services:

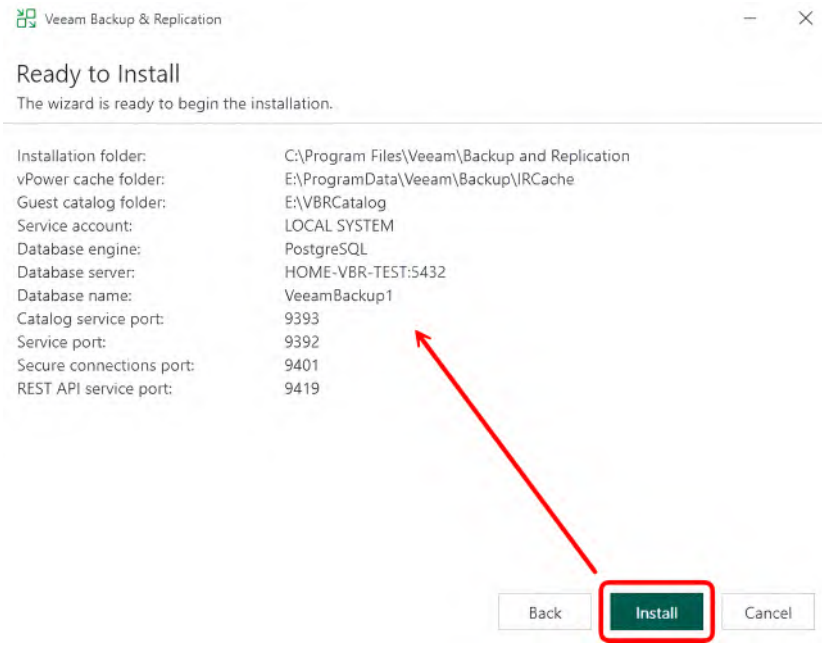


Figure 1.11: Ready to Install – checking for updates

11. Click **Install** to proceed with the installation and start setting up the components that work together with the backup server.

> **Note**
>
> PostgreSQL Server 17.6 gets installed during the installation, but as of this writing, the latest release is 18.1. Please check this link for the most recently supported version - `https://helpcenter.veeam.com/rn/veeam_backup_13_0_1_release_notes.html#system-requirements-backup-infrastructure-veeam-backup---replication-server-on-windows-configuration-database`

12. You will now see the installation progress dialog after clicking **Install**.

Technet24

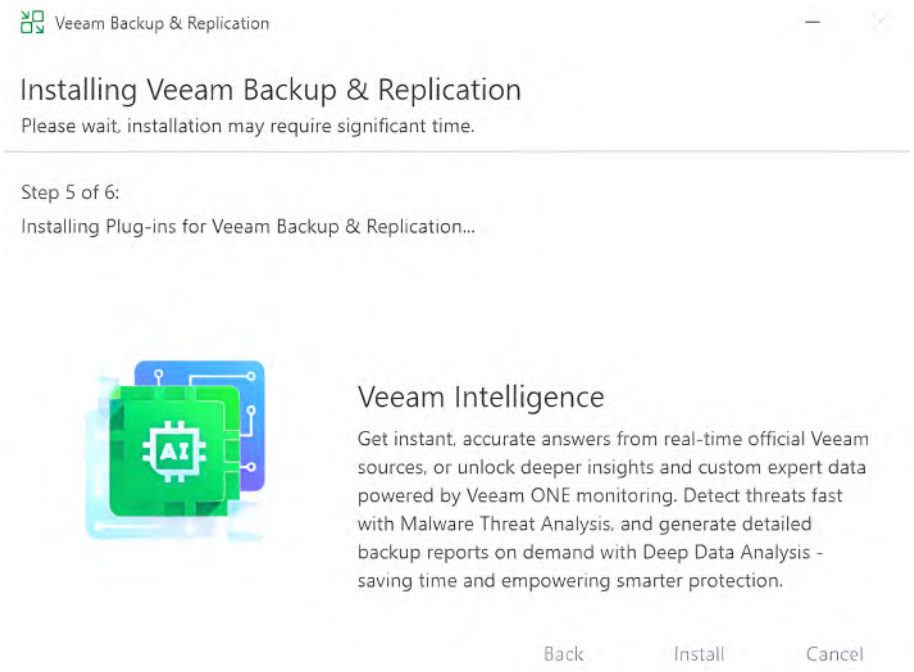


Figure 1.12: Installation progress

13. Once complete, you can click **FINISH** to complete the installation.

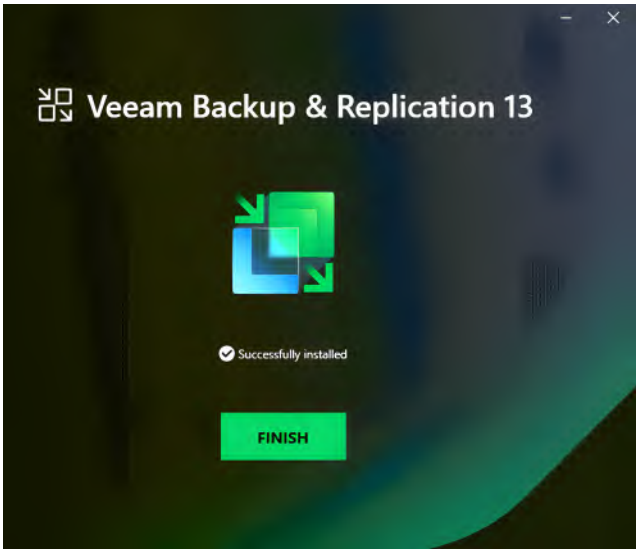


Figure 1.13: Installation complete

Now, let's start configuring the following required settings for Veeam to work with VMware:

- **Repository Server**: The server that is used for storing backup files.
- **Proxy Servers**: The servers that perform all the backup tasks.
- **VMware vCenter Credentials**: This is used to connect and see your clusters, hosts, vApps, and virtual machines. vCenter Server is not required, as standalone ESXi hosts are also supported, provided they are licensed in VMware.

Now, follow these steps to configure the required settings:

1. When you first launch the **Veeam Backup & Replication** console, you will be taken directly to the **Home** tab, and **Jobs** will be in focus.

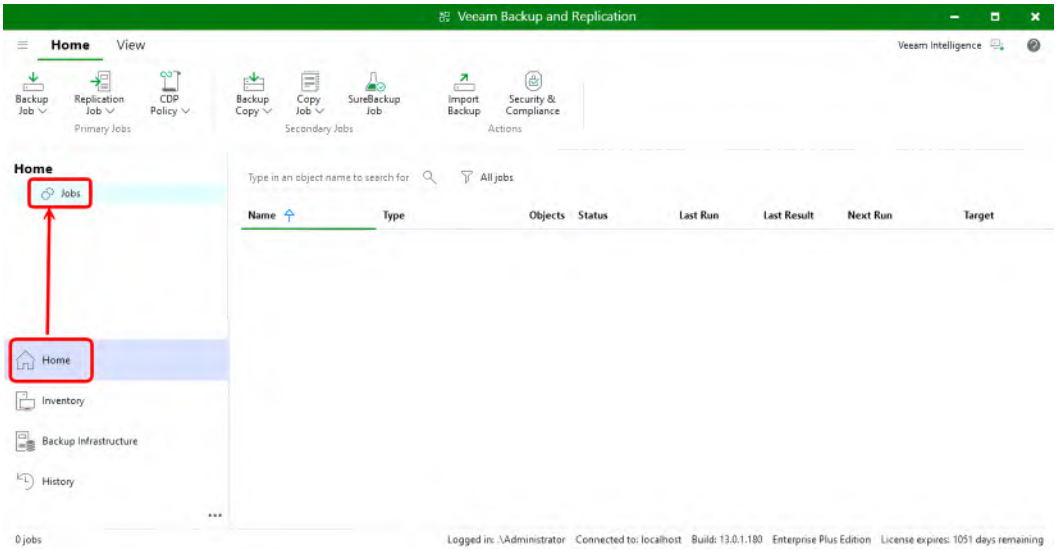


Figure 1.14: Initial console screen

2. Click on the **Backup Infrastructure** tab in the console to start adding servers.
3. This screen is where we will add the Virtual Center so that you can start backing up your virtual machines. Click on **Managed Servers** in the tree and then select the **Add Server** button in the toolbar to begin this process.

Technet24

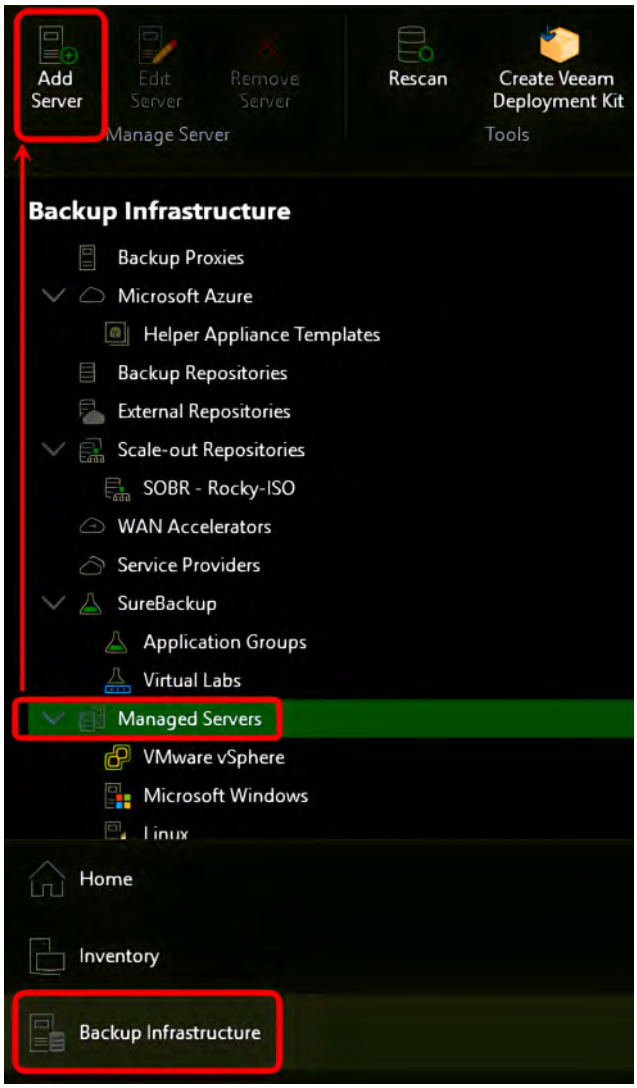


Figure 1.15: Managed Servers where you add your vSphere vCenter server

4. You will be prompted to select what kind of server to add. Choose **Virtualization Platform** and then **VMware vSphere**, where you will select either **vSphere** or **vCloud Director**.



Figure 1.16: vSphere or vCloud Director selection

You would typically select vSphere here; however, if you have vCloud Director in your environment, you can also choose this option. When you choose **vSphere**, you will be prompted with the following information to complete the connection:

- The DNS or IP address of your vCenter server (DNS is the preferred method because if the IP address of your vCenter server changes, then Veeam will still be able to connect)
- Credentials: This can either be a `vsphere.local` user or a domain account that is set up for access. You can also, at this point in the wizard, add credentials to the Veeam console using the **Add** button.

5. After selecting the required credentials or adding new ones, click **Apply** and **Finish** to complete the VMware vSphere setup. You will now see your vCenter server listed under the **Inventory** section of the console and can browse the hosts and virtual machines.

Now, let's look at the next piece required for the infrastructure: the **Proxy server**. By default, the **Veeam Backup & Replication** server serves as both your **VMware Backup Proxy** and **File Backup Proxy**. Due to my lab's limitations, I will use this server as an example; however, in the real world, you would typically add multiple Proxy servers to your environment for improved performance and adherence to best practices. Additionally, based on best practices, you would usually disable the Veeam Backup and Replication server from serving as a Proxy server, allowing the other Proxy servers to handle the workload.

Technet24

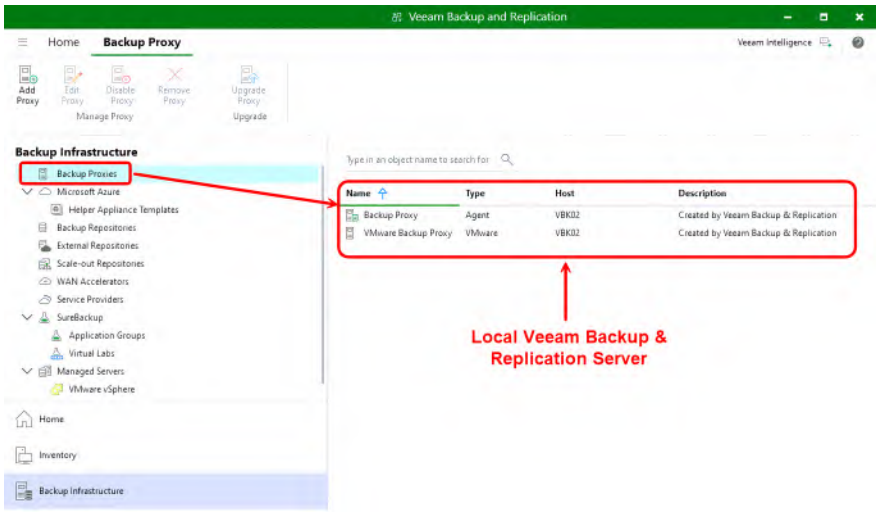


Figure 1.17: Default backup proxy is the Veeam server

The next component you will require is a **Repository Server**, which is the location where Veeam Backup and Replication will store your backup files. By default, Veeam Backup and Replication creates a **Default Backup Repository**, typically on the biggest drive attached to your backup server. This location will normally store **Configuration Backups**. As shown in *Figure 1.18*, there are multiple options for adding a repository.

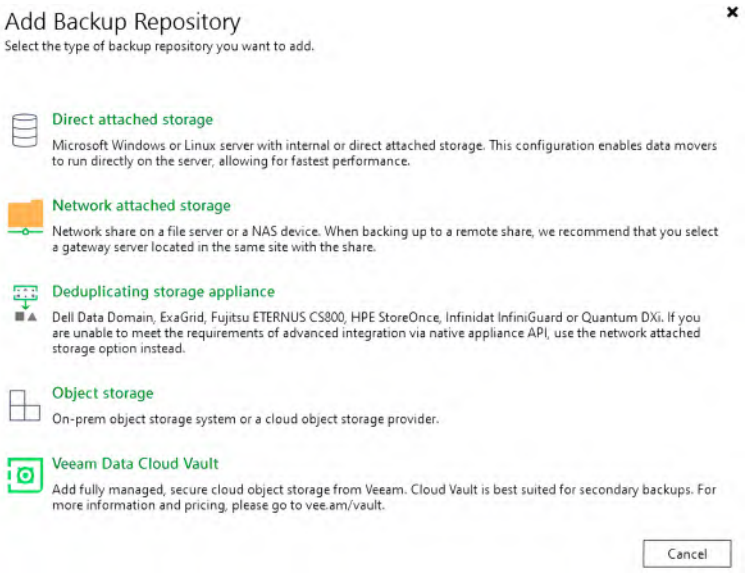


Figure 1.18: Add Backup Repository selection

The first three selections are for block storage. The last option is **Veeam Data Cloud Vault**, which can now be used for backups directly due to the enhancement in v13.

Figure 1.19 shows a backup job configured to go to a SOBR with Object storage as the Performance Tier.

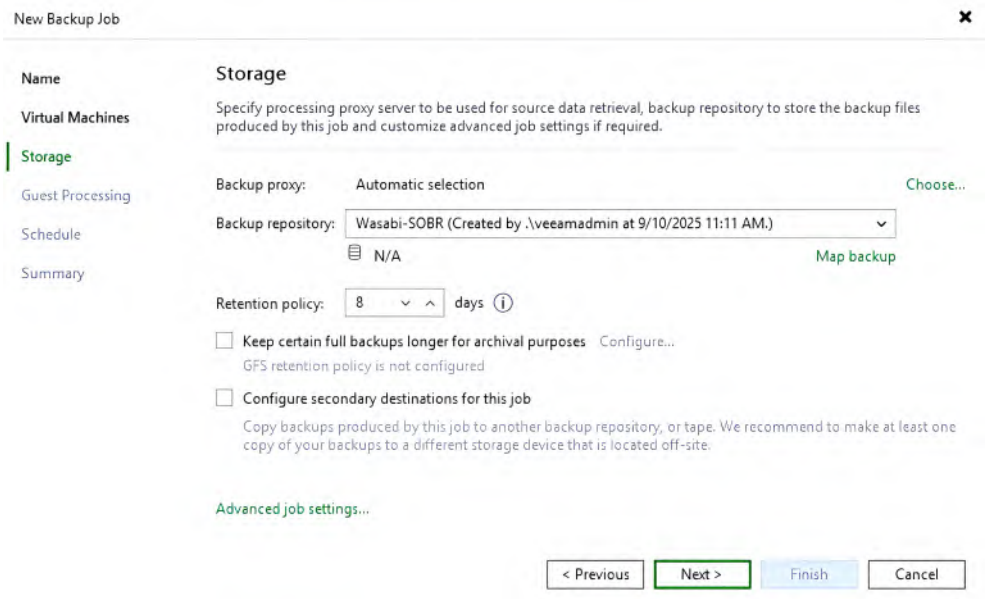


Figure 1.19: Backup job using SOBR with direct to Object Storage

Note

Offloading data happens with the Capacity and Archive tiers.

And *Figure 1.20* shows the SOBR configured with Wasabi Object Storage as the Performance Tier.

Technet24

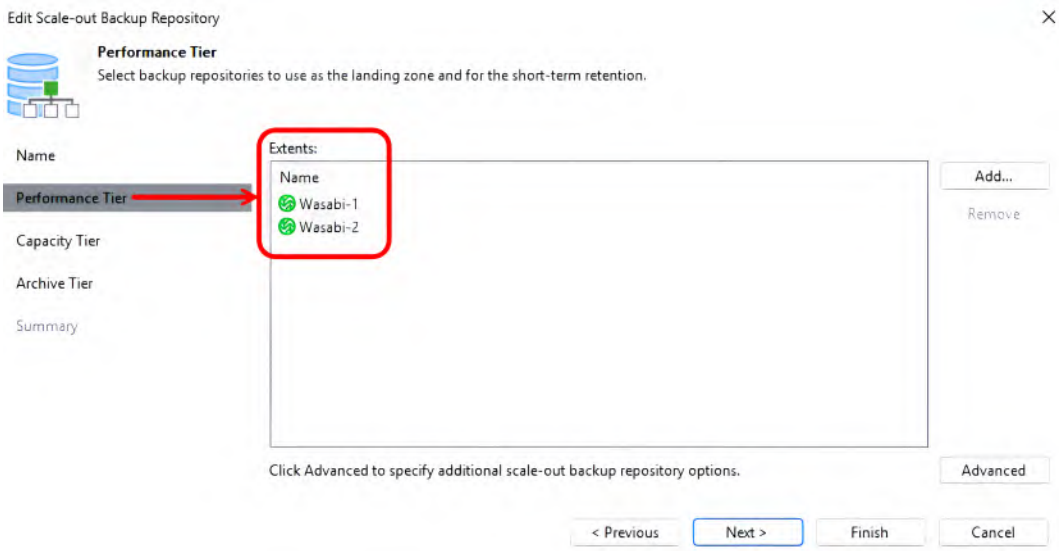


Figure 1.20: Scale-Out Repository using Wasabi Object Storage for Performance Tier

You have now installed and done the basic configuration required for **Veeam Backup & Replication**. We will now look at how to optimize **Proxy servers** and **Repository servers**.

Configuring and optimizing proxy servers

Proxy servers are the workhorses of the Veeam Backup and Replication v13 application, handling all the heavy lifting and processing tasks for backup and restore jobs. When you set up Veeam, ensure that the Proxy Servers are configured according to the best practices as outlined in the following Veeam webpages:

- `https://bp.veeam.com/vbr/2_Design_Structures/D_Veeam_Components/D_backup_proxies/vmware_proxies.html`
- `https://helpcenter.veeam.com/docs/vbr/userguide/proxies.html?ver=13`

Veeam Backup and Replication v13 continues to use Linux proxies, providing companies with an alternative to move away from licensing a Windows server. However, not all Linux versions are supported, as noted in `https://helpcenter.veeam.com/docs/vbr/userguide/system_requirements.html?ver=13#vmware-backup-proxy`

When you deploy a Proxy server, if it is a Windows system, Veeam Backup and Replication will install two components on the server:

- **Veeam Installer Service**: This is used to check the server and upgrade software as required.
- **Veeam Data Mover**: This is the processing engine for the Proxy server and does all the required tasks.

When you deploy a Linux Proxy server using the new Veeam Infrastructure Appliance ISO file, all necessary components are included, including the required services.

Veeam Backup and Replication Proxy Servers use a **transport mode** to retrieve data during backup. Three standard modes are available, and they are listed here in order, starting with the most efficient method:

- **Direct Storage Access**: The proxy is placed in the same network as your storage arrays and can retrieve data directly.

> **Tip**
>
> Note that when using Direct Storage Access (DSA), certain limitations apply. Refer to the Veeam documentation link - `https://helpcenter.veeam.com/docs/vbr/userguide/transport_modes.html?ver=13#transport-mode-limitations` – and review the note at the bottom of the page, "*These backup infrastructure components may support additional transport modes and block protocols if configured manually using root access permissions.*"

- **Virtual Appliance**: This mode mounts the VMDK files to the Proxy server for what we typically call **Hot-Add Mode** to back up the server data.
- **Network**: This mode is the least efficient but gets used when the previous methods are unavailable. It moves the data through your network stack. It is recommended not to use 1 GB but instead 10 GB. 10 GB is recommended to allow maximum throughput and performance when using the network mode.

In addition to these standard transport modes, which are provided natively for VMware environments, Veeam provides two other transport modes: **Backup from Storage Snapshots** and **Direct NFS**. These provide storage-specific transport options for NFS systems and storage systems that integrate with Veeam.

See the integration with storage systems guide for more details: `https://helpcenter.veeam.com/docs/vbr/userguide/storage_integration.html?ver=13`.

Technet24

Along with the transport modes, there are specific tasks that the Proxy server performs:

- Retrieving the VM data from storage
- Compressing
- Deduplicating
- Encrypting
- Sending the data to the backup Repository server (backup job) or another backup proxy server (replication job)

Veeam Proxy Servers leverage what is known as VMware **vStorage APIs – Data Protection** (formerly known as VMware vStorage APIs for Data Protection or VADP) when using all transport modes except Backup from Storage Snapshots and Direct NFS.

It would be best to keep the following considerations in mind while planning your Proxy servers:

- **Operating System**: Most software vendors will always recommend the latest and greatest, so if you choose Windows, choose 2025. Alternatively, you can select Linux using the latest release such as *Ubuntu 24.04.3 LTS*. For Linux VMware, backup proxies support all transport modes as of Veeam Backup & Replication v11a and support being a CDP Proxy in v12+.
- **Proxy placement**: Depending on the transport mode for the server, you will need to place it as close to the servers you want to back up, such as on a specific host in VMware. The closer to the source data, the better!
- **Proxy sizing**: This can be tricky to determine and will depend on the physical or virtual server. Veeam Proxy Servers complete **tasks** where one virtual disk is processed for a VM or one physical disk is processed for a server. Therefore, Veeam recommends one physical core or one vCPU and 2 GB of RAM per task.

Veeam has a formula to calculate the required resources for a Proxy server:

- **D** = Source data in MB
- **W** = Backup window in seconds
- **T** = Throughput in MB/s = D/W
- **CR** = Change rate
- **CF** = Cores required for full backup = T/100
- **CI** = Cores required for incremental backup = (T * CR)/25

Based on these requirements, we can use a data sample to perform the calculations:

- 500 virtual machines

- 100 TB of data
- An 8-hour backup window
- 10% change rate

Using these numbers, we can perform the following calculations:

$$D = 100\,TB * 1024 * 1024 = 104{,}857{,}600\,MB\ (Data\ is\ converted\ to\ MB)$$

$$W = 8\,hours * 3600\,seconds = 28{,}800\,seconds$$

$$T = 104{,}857{,}600 / 28{,}800 = 3{,}641\,MB/s$$

We can use the numbers we calculated to determine the required number of cores needed to run both full and incremental backups to meet our defined SLA:

$$CF = T / 100 \rightarrow CF\ (Full) = 3{,}641 / 100 \sim 36\,cores$$

$$CI = (T * CR) / 25 \rightarrow CI\ (Incremental) = (3{,}641 * 10) / 25 \sim 15\,cores$$

Based on our calculations and considering that you require 2 GB of RAM for each task, you need a virtual server with 36 vCPUs and 72 GB of RAM. This size may seem like a considerable server, but keep in mind that it uses sample data. Your calculations will likely be much smaller or more extensive, depending on your dataset.

If you plan to use a physical server as a Proxy, it should have 2–10 core CPUs. In the case of our sample data, two physical servers are required. If you are using virtual servers for proxies, the best practice is to configure them with a maximum of 8 vCPUs and add as many as needed for your environment. In this case, we would need five servers.

Tip

If you only want to size things based on incremental backups, your requirements are less than half of the full backup sizing – 15 vCPUs and 30 GB of RAM.

There are limitations to consider when using Proxy servers, particularly in terms of job processing and performance. As we noted previously, a Proxy server performs **tasks** which are assigned CPU resources. Concurrent task processing depends on the resources available in your infrastructure and the number of Proxy servers you have deployed.

As shown in *Figure 1.21*, when it comes to adding a Proxy server to Veeam Backup and Replication, there is the **Max concurrent tasks** option, which correlates to the number of CPUs that are assigned.

Technet24

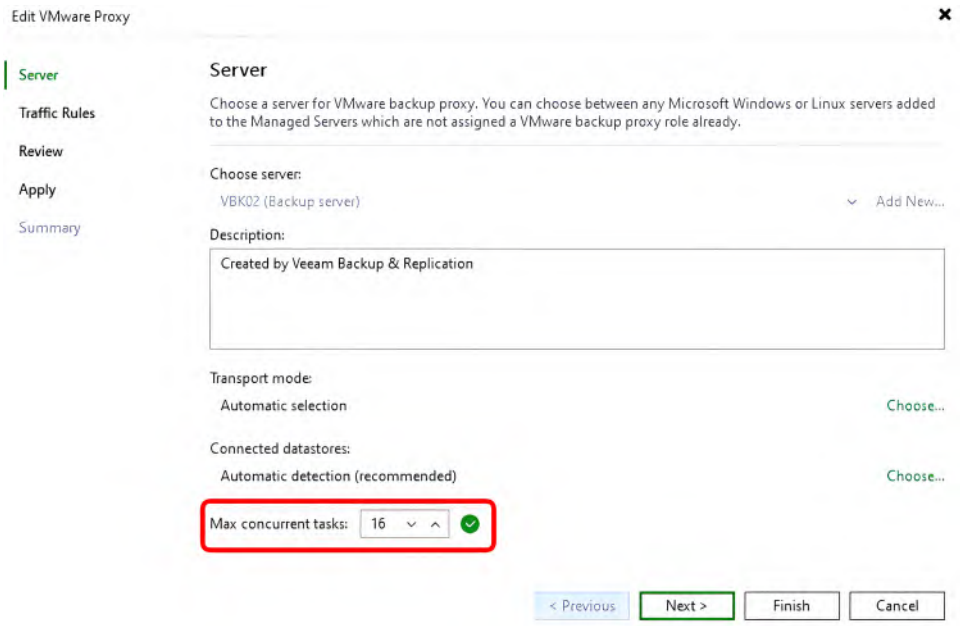


Figure 1.21: Max concurrent tasks limitation for Proxy

Task limits can be found at: `https://helpcenter.veeam.com/docs/vbr/userguide/limiting_tasks.html?ver=13`.

> **Note**
>
> The tasks of a Proxy server impact job performance. For example, if you had a Proxy server with 8 CPUs and added two virtual machines for backup, one with four disks and another with six disks, the Proxy server would process only eight of the ten disks in parallel. The remaining two disks would have to wait for resources before backing up in previous versions of Veeam.

With Veeam Backup and Replication v13, when you add a new Proxy server, it automatically sets the concurrent tasks based on the number of CPUs/vCPUs that the physical or virtual server has. In the case of my lab, the Veeam server has eight vCPU, so it was configured with sixteen concurrent tasks. Yes, this does not follow the sizing rule, but Veeam allows twice the number of CPUs/vCPUs for the concurrent task count.

You should now be able to right-size your Proxy Servers regarding CPU and RAM, understand proxy placement and how it processes tasks. Proxy Servers send data to Repository servers, which is the focus of the next section.

Setting up repository servers for success

A Repository server is a storage location for your backups, so setting it up correctly the first time will ensure you have optimal performance. When creating a repository, it is always a good idea to follow the Veeam Backup and Replication best practices - `https://bp.veeam.com/vbr/2_Design_Structures/D_Veeam_Components/D_backup_repositories/`.

The following are some considerations to keep in mind while setting up a repository:

- **ReFS/XFS**: With Windows 2019/2022/2025, ensure you format your repository drive(s) as ReFS with 64k block sizing to take advantage of space savings for Synthetic Fulls and GFS. For Linux, you need to set up XFS and Reflink to take advantage of space-saving and Fast Cloning. In both situations, storage efficiency will be realized for Synthetic Full backups. This efficiency prevents duplication but is not a form of deduplication.
- **Sizing**: Ensure that you adhere to the Veeam Backup & Replication recommendation of 1 core and 4 GB of RAM per repository task. Just like Proxy servers, your Repository servers have task limits as well. At a minimum, you need two cores and 8 GB of RAM.

When calculating the sizing requirements, consider your Proxy servers and the amount of CPU configured. Then, use a 3:1 ratio for the core count on a Repository server.

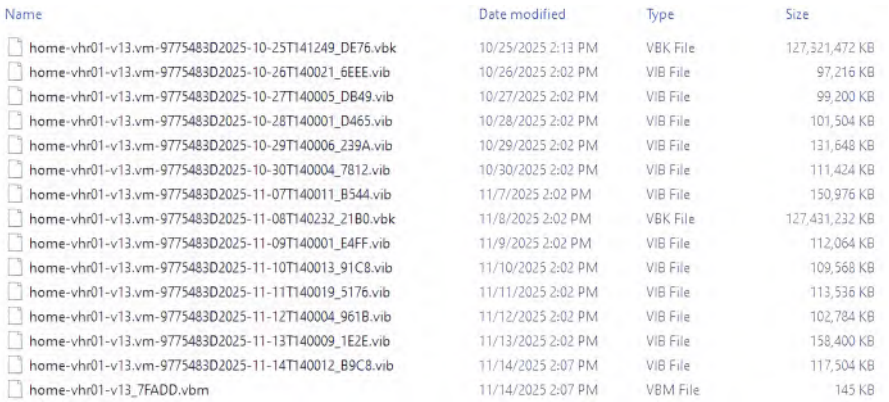
For example, if your Proxy server is configured with 8 CPUs, you need to configure the Repository server with 2 CPUs based on the 3:1 rule. To configure the RAM, you must multiply the CPU count by 4 to end up with 8 GB of RAM.

When you use the Windows ReFS file system as your repository, you must consider the overhead required for the file system and allocate an additional 0.5 GB of RAM per Terabyte of ReFS.

Setting up task limits for a Repository server differs from setting task limits for a Proxy server due to how tasks are consumed. The setting you choose will be handled differently:

- **Per-VM Backup Files**: Selected by default in Veeam Backup & Replication v13, this creates a backup chain per VM located in a job; rather than having a chain for all VMs together, they get separated, so each VM has its chain with a VBM, VBK, and VIB files. Therefore, if the backup job has 10 virtual machines, it will consume 10 repository tasks and 10 proxy tasks.
  Now an example of the new Per VM Backup chain from v13 is shown in *Figure 1.21* where each VM will have separate VBM, VBK, and VIB files.

Technet24



Name	Date modified	Type	Size
home-vhr01-v13.vm-9775483D2025-10-25T141249_DE76.vbk	10/25/2025 2:13 PM	VBK File	127,321,472 KB
home-vhr01-v13.vm-9775483D2025-10-26T140021_6EEE.vib	10/26/2025 2:02 PM	VIB File	97,216 KB
home-vhr01-v13.vm-9775483D2025-10-27T140005_DB49.vib	10/27/2025 2:02 PM	VIB File	99,200 KB
home-vhr01-v13.vm-9775483D2025-10-28T140001_D465.vib	10/28/2025 2:02 PM	VIB File	101,504 KB
home-vhr01-v13.vm-9775483D2025-10-29T140006_239A.vib	10/29/2025 2:02 PM	VIB File	131,648 KB
home-vhr01-v13.vm-9775483D2025-10-30T140004_7812.vib	10/30/2025 2:02 PM	VIB File	111,424 KB
home-vhr01-v13.vm-9775483D2025-11-07T140011_B544.vib	11/7/2025 2:02 PM	VIB File	150,976 KB
home-vhr01-v13.vm-9775483D2025-11-08T140232_21B0.vbk	11/8/2025 2:02 PM	VBK File	127,431,232 KB
home-vhr01-v13.vm-9775483D2025-11-09T140001_E4FF.vib	11/9/2025 2:02 PM	VIB File	112,064 KB
home-vhr01-v13.vm-9775483D2025-11-10T140013_91C8.vib	11/10/2025 2:02 PM	VIB File	109,568 KB
home-vhr01-v13.vm-9775483D2025-11-11T140019_5176.vib	11/11/2025 2:02 PM	VIB File	113,536 KB
home-vhr01-v13.vm-9775483D2025-11-12T140004_961B.vib	11/12/2025 2:02 PM	VIB File	102,784 KB
home-vhr01-v13.vm-9775483D2025-11-13T140009_1E2E.vib	11/13/2025 2:02 PM	VIB File	158,400 KB
home-vhr01-v13.vm-9775483D2025-11-14T140012_B9C8.vib	11/14/2025 2:07 PM	VIB File	117,504 KB
home-vhr01-v13_7FADD.vbm	11/14/2025 2:07 PM	VBM File	145 KB

Figure 1.22: Per VM Backup Chain in v13

- **No Per-VM Selection**: This feature for backup chains is now deprecated in Veeam Backup & Replication v13.

> **Note**
>
> To find out more about task limits, go to `https://helpcenter.veeam.com/docs/vbr/userguide/limiting_tasks.html?ver=13`.

When you're setting up a repository for the first time, you can set the task limit as follows:

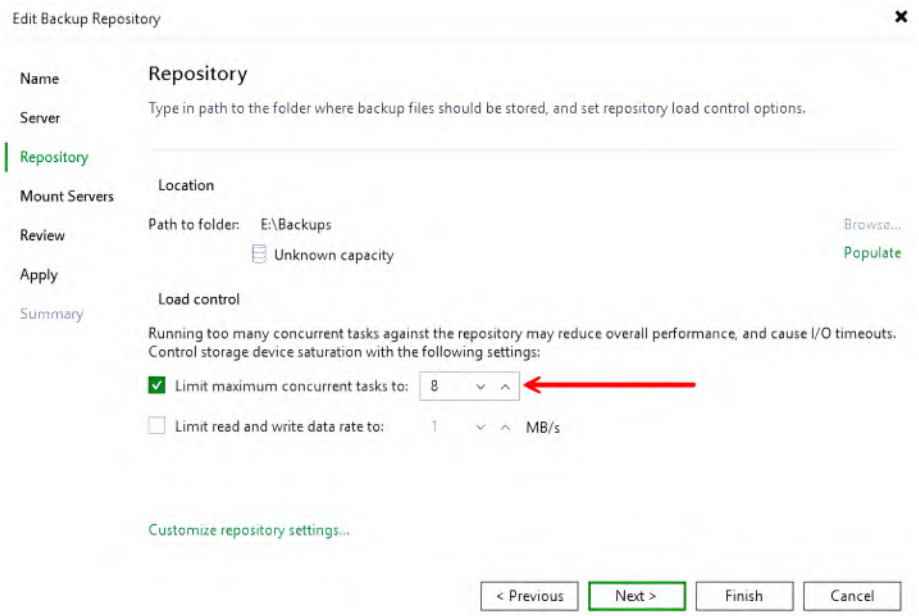


Figure 1.23: Repository task limit

Note

When you limit the number of tasks per repository and have jobs with many virtual machines requiring backups, this will be one of the bottlenecks in your environment. You also need to ensure that you do not set the limit too high, as this could overwhelm your storage and cause performance degradation. Make sure you test all your components and the resources available for your backup infrastructure.

After completing this section, you should be able to choose the type of file system you wish to use for your repository and size it correctly based on your CPU and RAM. We also discussed the per-VM method and how it gets created. We will utilize this knowledge to integrate it with the creation of a Scale-Out Backup Repository.

Understanding the Scale-Out backup repository

So, what is a **Scale-Out Backup Repository**, or **SOBR**, you ask? An SOBR uses multiple backup repositories, called Performance extents, to create a sizable horizontal scaling repository

Technet24

system. Veeam Backup and Replication can use multiple repositories of various types, such as the following:

- **Windows Backup Repositories**: NTFS or the recommended ReFS
- **Linux Backup Repositories**: XFS with Reflink
- **Object Storage**: Direct to Object Storage is available since v12 to use for all tiers of a SOBR, including the Performance Tier
- **Shared Folder**: **NFS**, **SMB**, or **CIFS**
- Deduplication storage appliances

SOBR can expand with on-premises storage, such as block storage or object storage, or a cloud-based object repository known as Capacity Extent. Veeam Backup and Replication combines the **Performance** extents, **Capacity** extents, and **Archive** extents into one to summarize their capacities.

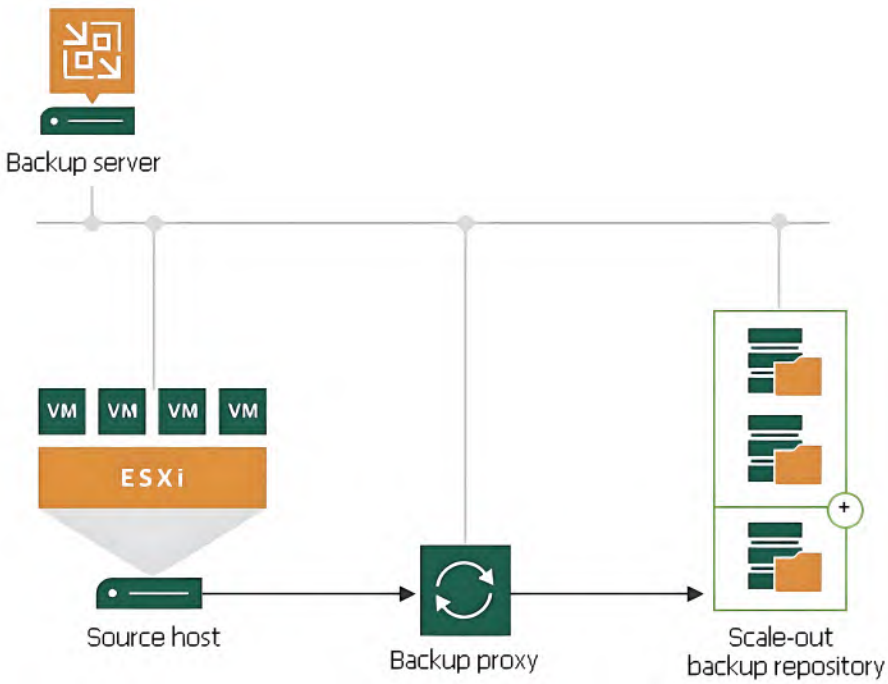


Figure 1.24: Scale-Out Backup Repository

The ability to use a SOBR is dependent on the license version that you are using with Veeam Backup & Replication:

- **Enterprise**: Allows for a total of two SOBRs with three active extents and one inactive extent.
- **Enterprise Plus & VUL (Veeam Universal License)**: Provides an unlimited number of SOBRs with as many performance and capacity extents needed, but only one Archive Tier per SOBR.

Tip

Should you downgrade your licensing from Enterprise Plus or Enterprise to Standard, you will lose the ability to target your jobs to the SOBR. You can, however, restore data from the SOBR.

The different license types limit the number of SOBRs and Extents per SOBR that you can configure. As we noted previously, there is a limit of two for **Enterprise** and an unlimited number for **Enterprise Plus/VUL**.

Tip

For optimal performance and manageability, it is best to keep your SOBR limited to three to four extents, if possible. If you are using object storage, then one of the components of the SOBR will be the Capacity Tier.

The scale-out repository works with many types of jobs or tasks in Veeam Backup and Replication:

- Backup jobs
- Backup copy jobs
- VeeamZIP jobs
- Agent Backups – Linux or Windows agent v2.0 or later
- NAS Backup jobs
- Nutanix AHV backup jobs
- Veeam Agent for MAC
- Veeam backups for Amazon and Microsoft Azure (via backup copy jobs)

Technet24

The next consideration to keep in mind is the limitations of using a SOBR, as there are certain things you cannot do:

- Only Enterprise, Enterprise Plus, and VUL licenses.
- You cannot use it as a target for configuration backup jobs, replication jobs, VM copy jobs, Veeam Agents v1.5 or earlier for Windows, or v1.0 Update 1 or earlier for Linux
- Adding a repository as an extent to an SOBR will not be allowed if there is an unsupported job using the repository
- Rotating drives are not supported
- You are unable to use the same extent in two scale-out repositories

Please refer to the following limitations page on the Veeam Backup and Replication website for more details - `https://helpcenter.veeam.com/docs/vbr/userguide/sobr_limitations.html?ver=13`.

When it comes to the makeup of the Scale-Out repository, there are three tiers:

- **Performance Tier**: Fast storage and fastest access to data can now include Object Storage, whether on-premise or cloud-based.
- **Capacity Tier**: Typically, this is object storage for archival and offloading capabilities
- **Archive Tier**: Additional object storage for long-term archival and infrequently accessed

Note

When adding Object Storage as the Performance Tier, ensure that each added instance is from the same vendor. For example, if the first Performance Tier you add is from Wasabi, the subsequent extents must also be Wasabi-based, as you cannot add something like Amazon S3. This will be crucial in planning if you decide to start using Object Storage in a SOBR.

The performance tier should use the fastest storage, whether block or object storage, to ensure the quickest possible access to files and restores. When you create a standard repository before adding it to a SOBR, there are specific settings retained in the SOBR:

- The number of simultaneous tasks it can perform
- The storage Read and Write speeds
- Data decompression settings on the storage
- The block alignment settings of the storage

The SOBR will not inherit a repository backed by rotating drives. Alternatively, if you choose to use the Per-VM backup option, it is enabled by default in a SOBR.

Note

If you decide to use cloud-based Object storage, keep in mind that it will not be the same as using block storage or on-premises object storage, as it will need to pull data from the cloud back to on-premises, which could slow down restores.

Another consideration is the backup file placement policy you will use. There are pros and cons to both, and specific operating systems, such as ReFS and XFS, require one over the other. The two types of placement policies are **Data locality** and **Performance**.

Please refer to the *Performance Tier* page on the Veeam Backup and Replication website for more information - `https://helpcenter.veeam.com/docs/vbr/userguide/backup_repository_sobr_extents.html?ver=13`

Data locality allows the SOBR to place all backup files within the same backup chain on a single extent, keeping related files together. The metadata file (VBM) is located on all extents in the SOBR for consistency and to ensure that files can be moved to different extents if needed. In contrast, the Performance policy will enable you to choose which extents to use for both full backup files (VBK) and incremental files (VIB).

For further information on backup placement, see this Veeam Backup and Replication page - `https://helpcenter.veeam.com/docs/vbr/userguide/backup_repository_sobr_placement.html?ver=13`

Now, when it comes to the Capacity Tier, there can be unlimited per scale-out, and it is required to be one of the following:

Technet24

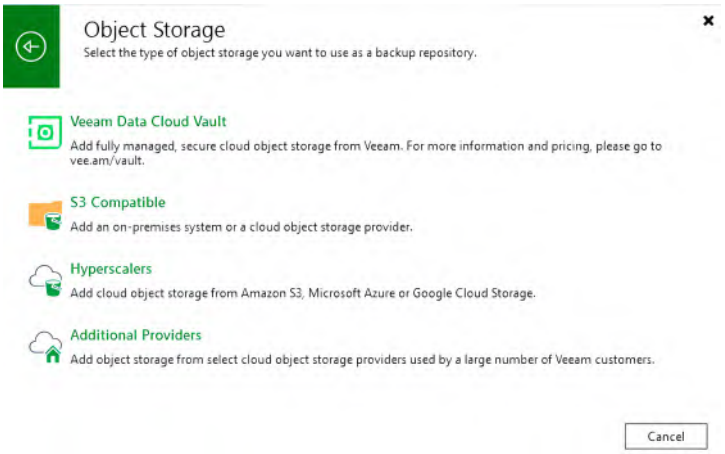


Figure 1.25: Object Storage options for Capacity Extent

Note

These same Object Storage options for the Capacity Tier now apply to the Performance Tier, and as noted, must be the same for each extent you want to add to the Performance Tier.

Using a Capacity Tier as part of your SOBR is suitable for the following reasons:

- You can tier off older data or when your SOBR reaches a specific percentage capacity to allow you to free up storage space.
- Company policy stipulates you keep a certain amount of data onsite, then all older data after X days is tiered off to the Capacity Tier.
- Using it falls into the 3-2-1-1-0 rule, where one copy of the information is offsite. See this blog post for more details on the 3-2-1-1-0 rule – `https://www.veeam.com/blog/321-backup-rule.html`.

You specify the Capacity tier after creating it as a standard repository and during the SOBR wizard.

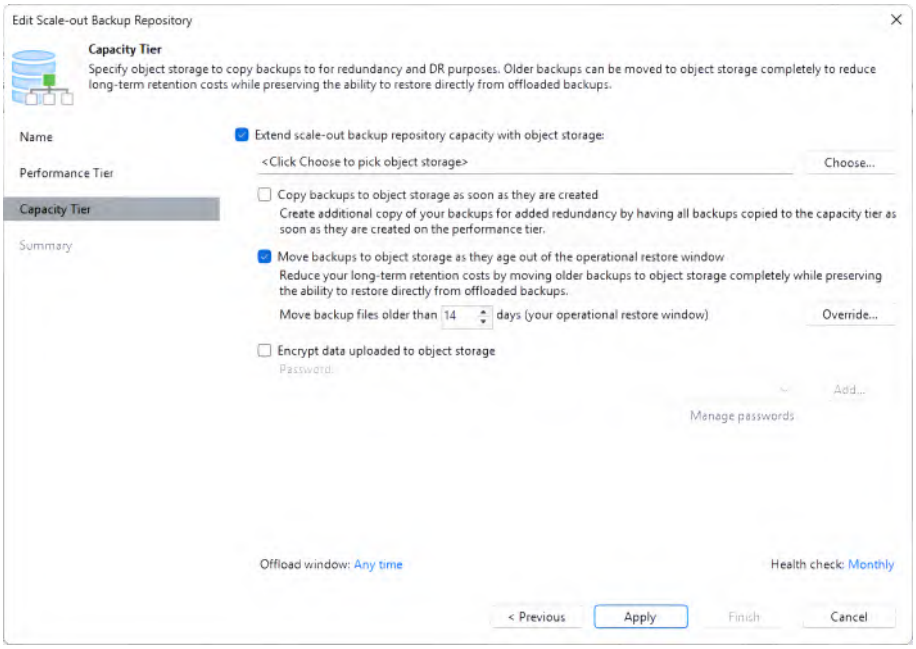


Figure 1.26: Capacity Tier of the Scale-Out wizard

Please visit the *Capacity Tier* page on the Veeam Backup and Replication website for more information: `https://helpcenter.veeam.com/docs/vbr/userguide/capacity_tier.html?ver=13`.

We will now tie everything we've learned together and create a SOBR.

1. First, open the Veeam Backup & Replication console and select the **Backup Infrastructure** section located at the bottom left. Then, click on the **Scale-out Repositories** option on the left.

Technet24

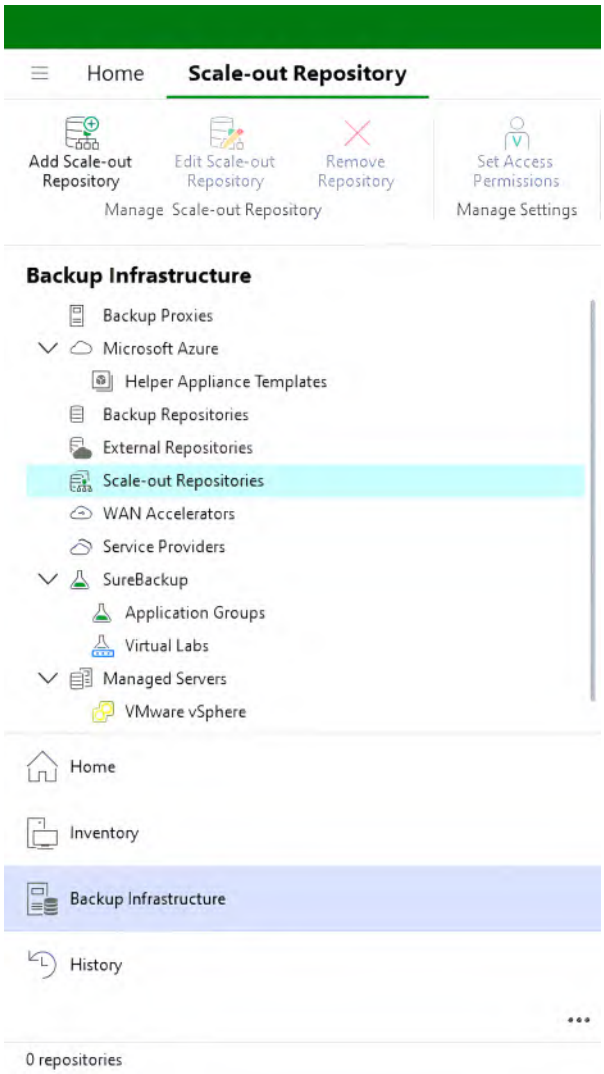


Figure 1.27: The Scale-out Repositories section of the console

2. Once in this section, you can either click the **Add Scale-out Repository** button in the toolbar or, in the right-hand pane, right-click and select **Add Scale-out Backup Repository**.
3. At this point, you must name the Scale-Out and give it a thoughtful description; the default name is **Scale-Out Backup Repository 1**. Then, click **Next** to go to the **Performance Tier** section of the wizard.

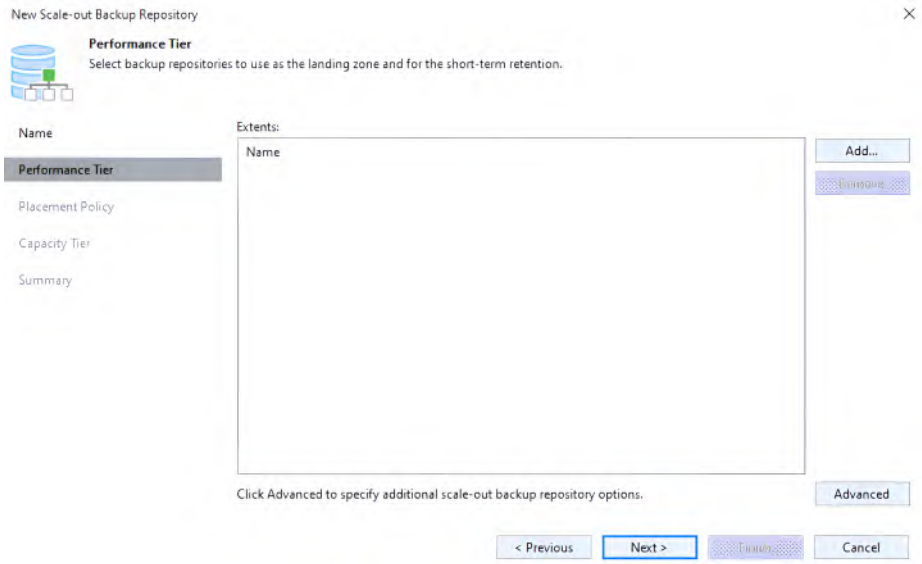


Figure 1.28: Scale-out wizard – Performance Tier

4. In this section, click the **Add...** button and choose the standard repositories that will be part of your Scale-out. You can also click the **Advanced** button to choose one option: **Perform full backup when the required extent is offline**.

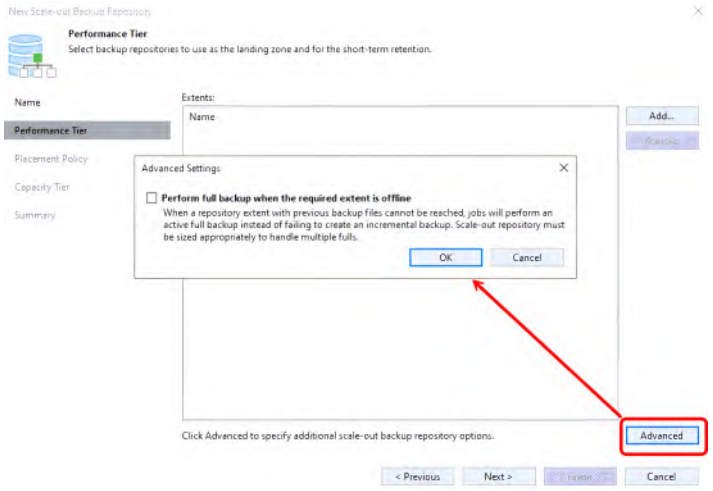


Figure 1.29: Advanced settings for the Performance Tier of SOBR

5. Click the **Next** button to proceed. At this point, you must pick your placement policy: **Data Locality** or **Performance**.

Technet24

As we mentioned, if you're using ReFS or XFS, you must select **Data Locality** to take advantage of each operating system's storage savings. Click **Next** after making your choice.

6. You can now choose to use **Capacity Tier** for your SOBR or click the **Apply** button to complete the process. Note that when you select a Capacity Tier, there are several options you can enable:
    - Copy backups to object storage as soon as they are created in the Performance Tier.
    - Move backups to object storage as they age out of the restore window. The default is 14 days. You can also click the **Override** button to specify when offloading should begin based on the space remaining below a certain percentage.
    - You can also encrypt your data upload to the object storage as another level of security:

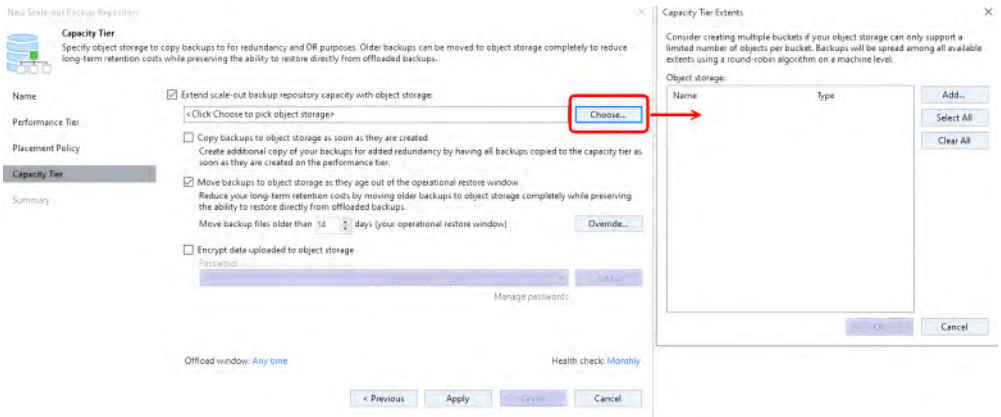


Figure 1.30: Capacity Tier selection for Scale-out

Note that some Capacity Tier targets support immutability. This feature is a crucial attribute in the fight against ransomware.

Please see the Veeam Readiness program to determine if your object storage is supported here: `https://www.veeam.com/solutions/alliance-partner/integrations-qualifications.html?type%5B%5D=immutability`

Once complete, you will see your new Scale-out repository. When you select it, you will see the performance tier extents and the Capacity Tier if you chose it.

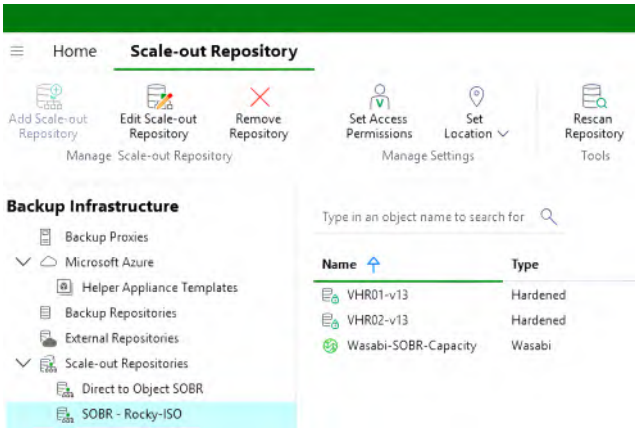


Figure 1.31: Scale-out repository created

For further information on the SOBR, visit this Veeam Backup and Replication page -`https://helpcenter.veeam.com/docs/vbr/userguide/sobr_add_console.html?ver=13`.

The final aspect to discuss is how to manage the SOBR after it has been created. Once created, you may need to do any of the following:

- Edit its settings to change the performance policy, for example
- You may need to re-scan the repository to update the configuration in the database.
- Extend the performance tier by adding another extent to the SOBR.
- Put an extent in maintenance mode to perform maintenance on the server that holds it or evacuate the backups to remove the extent.
- Switch an extent into **Sealed Mode**, where you do not want any more writes to it but can still restore from it. This process allows you to replace the extent with a new one.
- Run a report on SOBR.
- Remove an extent from the SOBR, which requires maintenance mode, evacuation, and removal.
- Remove the SOBR altogether.

Let's now briefly discuss the two additional features in Veeam Backup and Replication v13 for managing SOBR and repositories in general: the **VeeaMover** and **SOBR Rebalance**.

VeeaMover

It is a built-in utility that simplifies operations for moving data between repositories and VMs between jobs. It includes the ability to move from ReFS to ReFS or XFS, retaining block clone savings (a similar feature to XFS as well). You can also move from NTFS to ReFS to take

Technet24

advantage of the block clone savings in ReFS. Additionally, VeeaMover facilitates moving VMs from one job to another.

SOBR rebalance

When you use a SOBR with multiple extents, depending on your placement policy, some of the extents fill up more than others. There is now a feature called SOBR Rebalance, which helps with this by spreading backups across the extents. You activate this option by pressing the **Ctrl** key and right-clicking on the SOBR name to see the function. This operation requires an outage; however, moving the files around, depending on their size, could take some time.

Once set up within Veeam Backup and Replication v13, the SOBR is self-sufficient. Still, there is maintenance that needs to be performed to ensure optimal performance and ample storage is available for backups.

For more information on SOBR management, please visit the following page on the `Veeam Backup and Replication website: https://helpcenter.`veeam.com/docs/vbr/userguide/managing_sobr_data.html?ver=13.

Upgrading Veeam backup and replication to v13

This section outlines the process for upgrading your existing **Veeam Backup & Replication** environment to version **13**. If you have Veeam Backup and Replication v12 installed, you can proceed with the following steps to conduct the upgrade.

Once your server is ready and you have downloaded the *ISO file* and mounted it, follow these steps to upgrade your server and components:

1. Run the `setup.exe` file on the mounted ISO drive

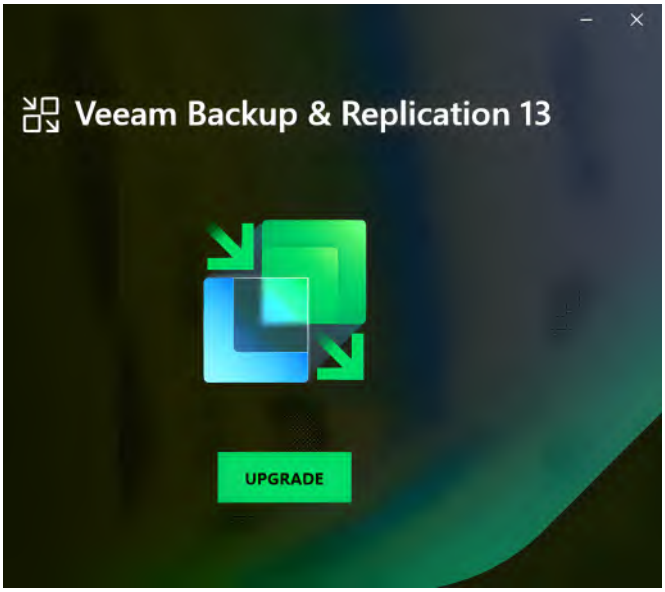


Figure 1.32: Main installation screen – Upgrade

2. Click on the **UPGRADE** button to proceed
3. You are then prompted to click on the component to upgrade, and in this case, **Veeam Backup & Replication**

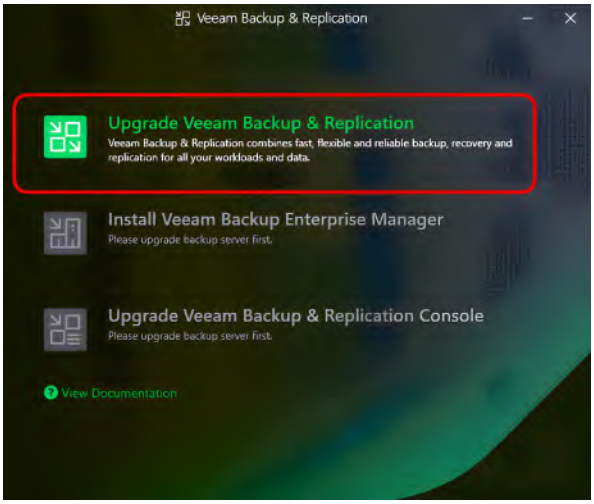


Figure 1.33: Upgrade selection

Technet24

> **Note**
>
> If Veeam Backup Enterprise Manager were installed on your server, you would be prompted to upgrade it before you could upgrade Veeam Backup and Replication.

4. After clicking **Upgrade Veeam Backup & Replication**, the upgrade installation will proceed
5. The **License Agreement** is the following dialog, where you must click **I Accept** to proceed

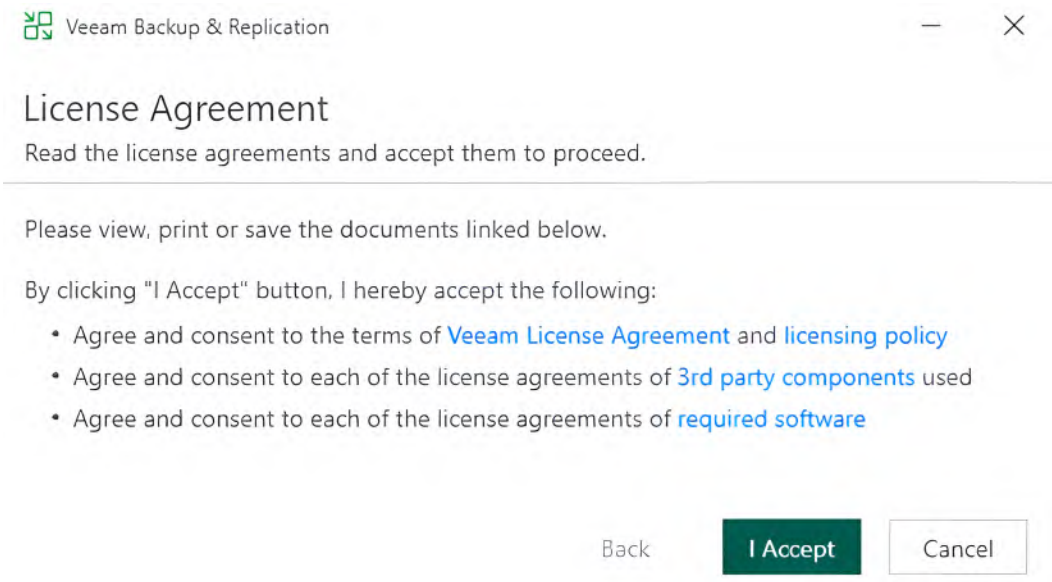


Figure 1.34: License Agreement dialog

6. The next screen is the License dialog, where you use the **Browse to local license file** link to enter your license. You can sign in to your Veeam account to select a license to install as well. Then, click **Next** to proceed.

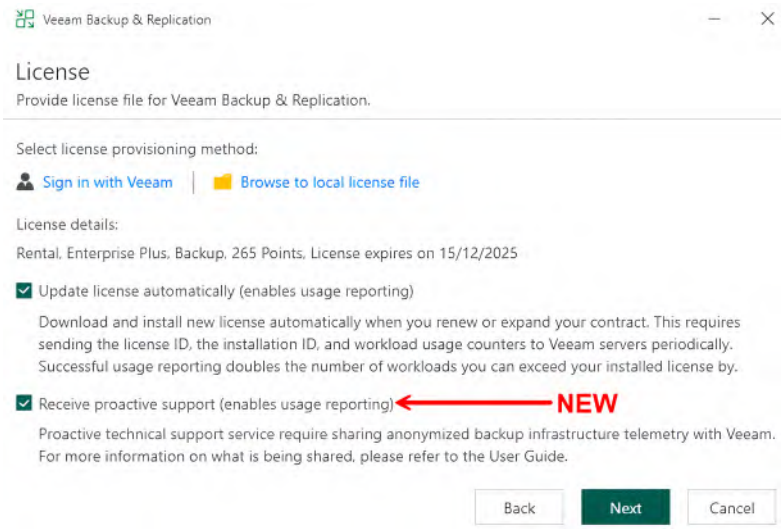


Figure 1.35: License selection dialog

Note

In the new Backup and Replication v13, a new option is available on the License dialog screen, allowing you to enable the **Receive proactive support** option.

7. The installer now does a system check and will enable any missing features required for the installation to continue

Technet24

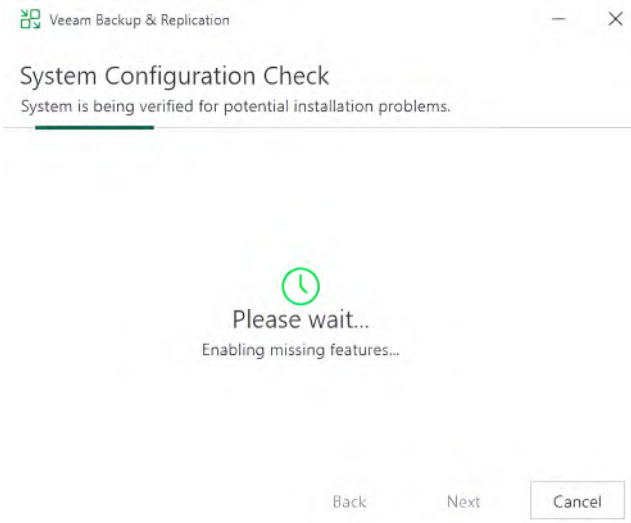


Figure 1.36: System Configuration Check

8. A **Configuration Check** is now performed to ensure there are no incompatibilities for the upgrade and shows you if there are so they can be resolved before proceeding with the upgrade. Click **Next** to proceed.

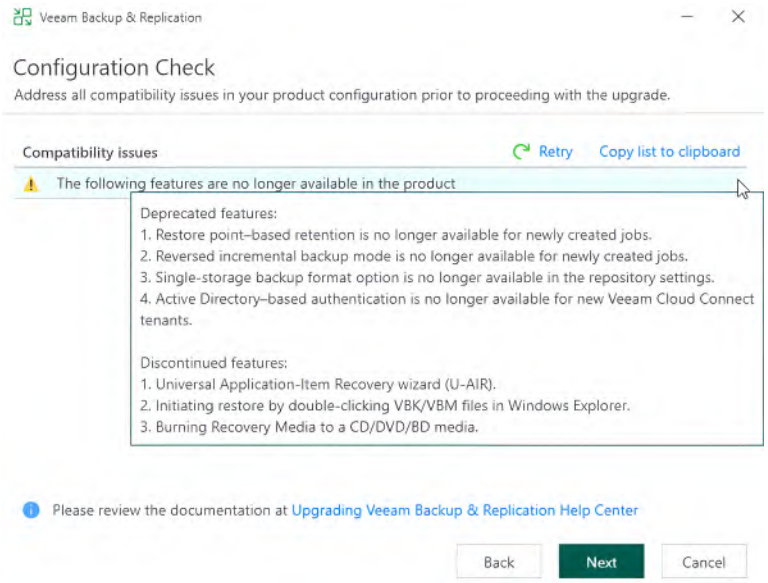


Figure 1.37: Configuration Check for Incompatibilities

> **Note**
>
> In my case, there are some features noted as deprecated in this release of version 13, as well as some discontinued features.

9. You are presented with the **Upgrade** dialog showing the components to be upgraded with the current and new versions. You can also check the **Update remote components automatically** option so that they are upgraded on console launch. Click the **Upgrade** button to proceed.

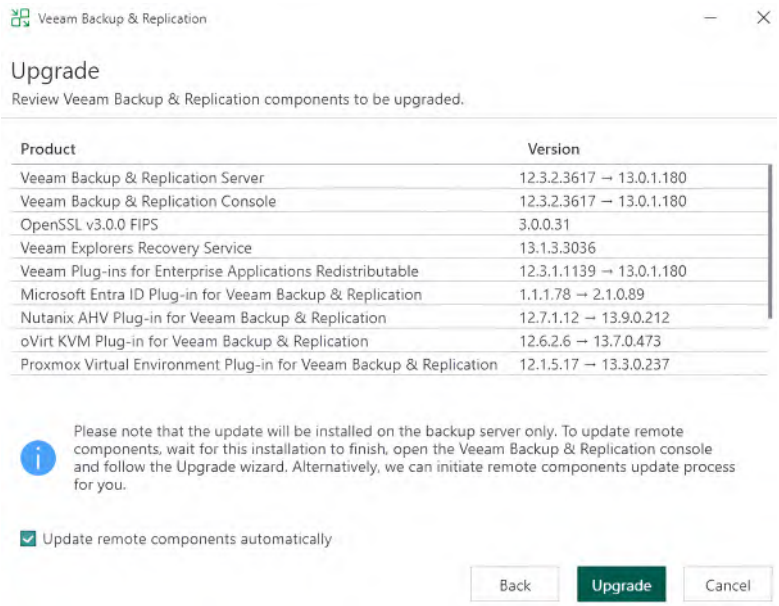


Figure 1.38: Ready to Upgrade dialog

10. After the progress dialog finishes, the upgrade is complete, and you can click **FINISH**.

Technet24

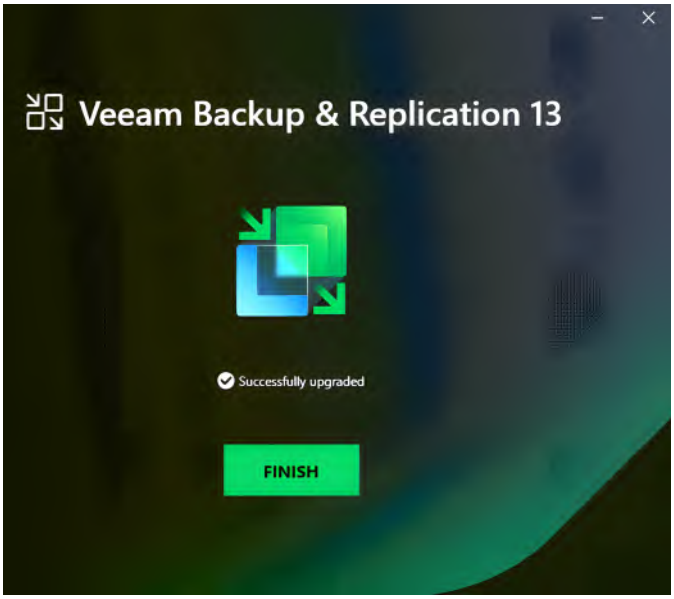


Figure 1.39: Upgrade complete

The upgrade process is now complete, and your jobs should continue to run as scheduled.

> **Note**
>
> Backup and Replication v13 now includes a software appliance – Veeam Software Appliance (VSA). It can be deployed via an ISO or the new OVA format. This will be covered in *Chapter 2, Veeam Software Appliance (VSA)*.

We have now looked at how you can upgrade your Veeam installation, which concludes this section of the chapter.

Summary

This chapter has provided you with the tools required to install Veeam Backup and Replication v13, as well as the components that comprise the installation. We discussed the prerequisites, including PostgreSQL, which is the default database that comes with the new release. Then, we addressed how to set up Proxy Servers, best practices for configuration, and optimal settings. This part was then followed by a discussion on repositories and how to create them and included best practices and optimizations for best performance. Afterward, we examined Scale-Out Backup Repositories and how to configure them, including the Performance Tier and Capacity Tier, as well as managing them once they are set up. Lastly, we examined the upgrade

process for Veeam Backup and Replication to update an existing instance from version 12 to version 13.

This chapter has ensured that you have covered all the basics. In the next chapter, we will cover the new Veeam Software Appliance (VSA) for Veeam Backup and Replication v13.

Further reading

- PostgreSQL Server - `https://www.postgresql.org/`
- Object Storage Repository in Veeam - `https://helpcenter.veeam.com/docs/vbr/userguide/object_storage_repository.html?ver=13`
- Backup Proxies - `https://helpcenter.veeam.com/docs/vbr/userguide/proxies.html?ver=13`

Technet24

2

Veeam Software Appliance (VSA)

This chapter will discuss the new Veeam Backup and Replication 13 appliance, known as the **Veeam Software Appliance** (**VSA**). You will learn about the latest option for deploying the VSA and how it is security-hardened right out of the box. We will also look at the new **High Availability** (**HA**) feature built into the appliance for disaster recovery. We will explore the new WebUI for management and the enhanced Role-Based Access Controls (RBAC), including the **Veeam Security Officer** (**veeamso**). Lastly, we will look at new features, such as auto-updates for all components, which Veeam now provides.

In this chapter, we're going to cover the following main topics:

- Introducing the Veeam software appliance
- Security hardening and high availability
- Web-based management interface and enhanced RBAC
- Centralized update management and auto updates

Technical requirements

One of the main requirements before you proceed with this chapter will be familiarity with VMware and with deploying **Open Virtualization Format** (**OVF**) or **Optical Disc Image** (**ISO**) templates. You will need the **Open Virtual Appliance** (**OVA**) or ISO file for VSA to complete this chapter.

Technet24

Introducing the Veeam software appliance

With the release of Veeam Backup and Replication v13, another deployment option has been introduced: **Veeam Software Appliance (VSA)**. All previous Veeam Backup and Replication versions were installed using the standard ISO file, whether for a new complete installation or an upgrade from an earlier version. With v13, you can now use an OVA or ISO file to deploy the VSA, a fully hardened Rocky Linux appliance provided by Veeam.

Those reading and hearing VSA for the first time might ask what it is and why Veeam chose to come out with an appliance. These are the main reasons:

- Both customers and partners asked for it
- The operating system is pre-hardened, updated, and tested by Veeam
- Simplified deployment using an all-in-one appliance
- Using your preferred hardware for deployment

The new VSA can be deployed via the usual ISO method on a virtual or physical server. With version 13, there is also an OVA option, a VMware template for deployment. The OVA file deployment method uses VMware's built-in mechanism to select the file and then create the virtual machine. Once you power on the virtual machine, it installs all required components and then walks you through the initial deployment wizard.

Note

When using the OVA deployment method, the two disk drives are set to 240GB each, which may not be suitable for larger environments. It is recommended to create a new virtual machine with the required specifications, then mount the ISO within the virtual machine to complete the installation. This allows you to use larger disks and more space for components like logs.

Please see the *Further reading* section for the link to the new appliance requirements.

You can download and deploy the VSA using an OVA file or ISO file from within the VMware environment, as shown in *Figure 2.1*.

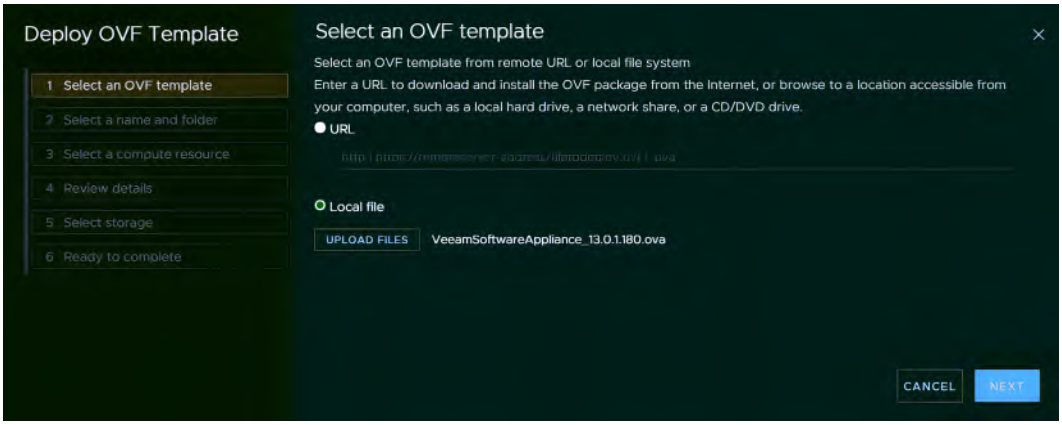


Figure 2.1: OVA deployment method in VMware

If you are booting from an ISO file you will see a menu with options to install Veeam Backup and Replication or Veeam Backup Enterprise Manager as seen in *Figure 2.2*.

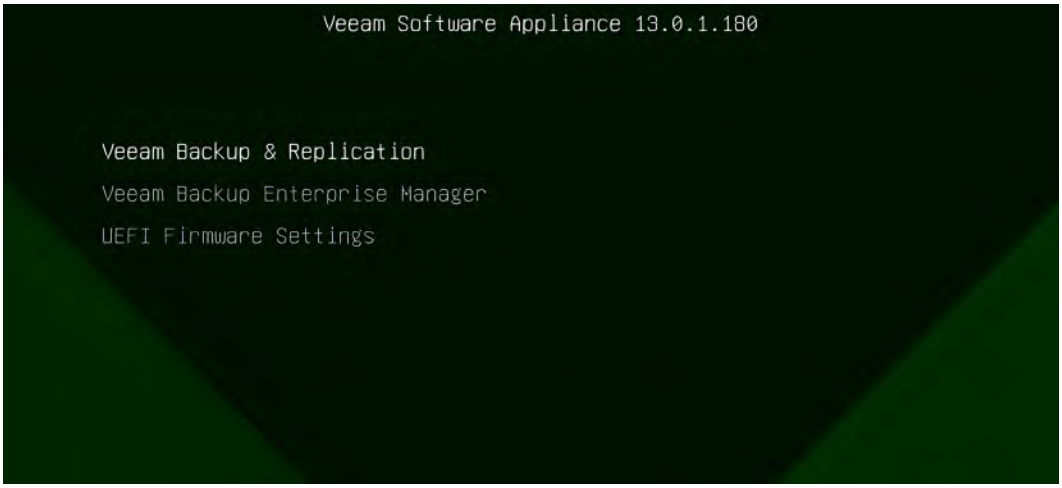


Figure 2.2: ISO deployment method

Here, you will see the different deployment options:

- **Veeam Backup & Replication**: This will install Veeam Backup & Replication v13 on the virtual machine or physical server

Technet24

- **Veeam Backup Enterprise Manager**: This will install Veeam Backup Enterprise Manager v13 on the virtual machine or physical server
- **UEFI Firmware Settings** :This takes you into the BIOS to check the UEFI settings for the system

> **Note**
>
> Veeam Backup and Replication v13 appliance includes PostgreSQL Server 17.6 locally, but the latest version is 18.1 as of this writing.

Installing Veeam backup and replication 13 with OVA and ISO

Installing Veeam Backup and Replication 13 is straightforward, with new OVA and ISO methods for deploying the software appliance. The Windows upgrade remains the same as noted in *Chapter 1, Installing and Upgrading Veeam*.

Follow these steps to install Veeam Backup and Replication:

1. After selecting **Veeam Backup & Replication** as shown in *Figure 2.2*, you then choose **Install** or **Reinstall**

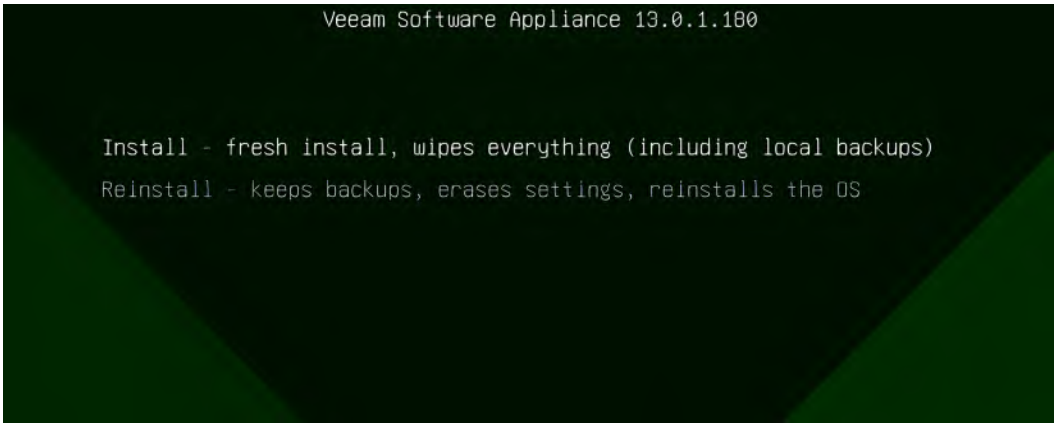


Figure 2.3: Install or Reinstall options

   Typically, you would select **Install** if this is your first deployment.

2. Once the system boots up, you will be presented with only one dialog that indicates all data will be wiped. Click **Yes**.

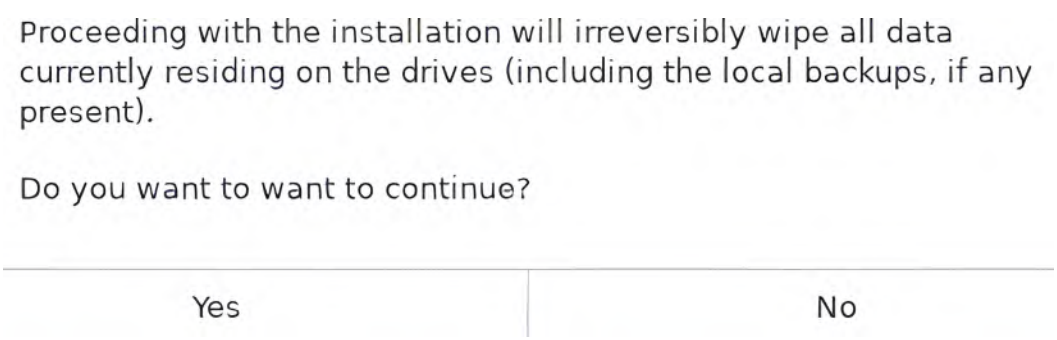


Figure 2.4: Proceed with install dialog

3. After you click **Yes**, the installation begins, and the Veeam installer gets the system ready for the software installation. You will configure the system after the installation is completed and the system reboots.

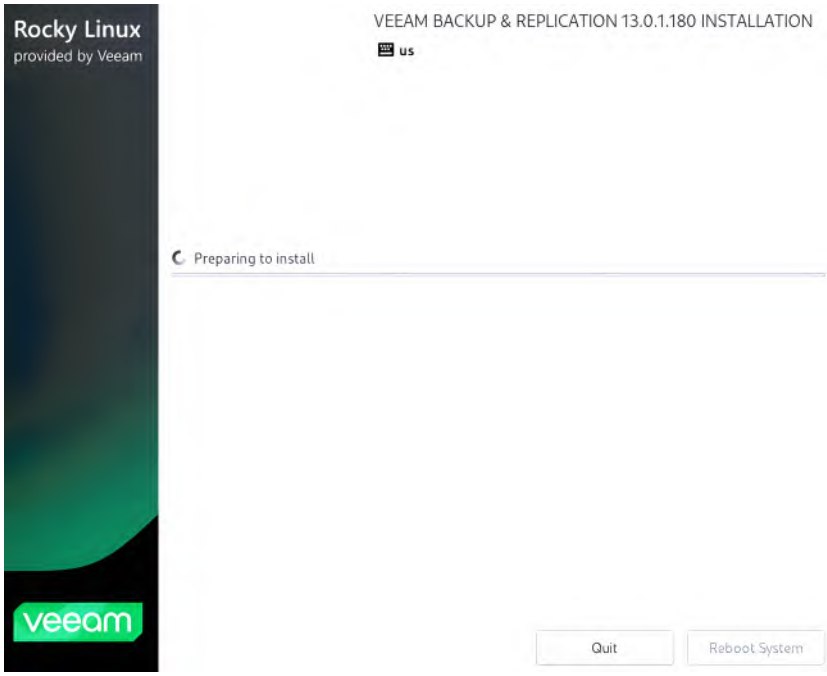


Figure 2.5: Installation of Veeam Backup and Replication software appliance based on Rocky

Technet24

> **Note**
>
> If you have used the OVA instead of the ISO, once deployed in VMware and powered on, there is no installation required. Once the system boots, you will be prompted to run the Initial Configuration Wizard, as with the ISO installation method above.

Figure 2.6 shows the **Initial Configuration Wizard**, which you need to complete to finalize the appliance configuration. Keep in mind that for the "veeamadmin" user, MFA is a requirement when setting this up. There is also a new role, **Security Officer**, in the latest appliance that you must set up in the Initial Configuration Wizard.

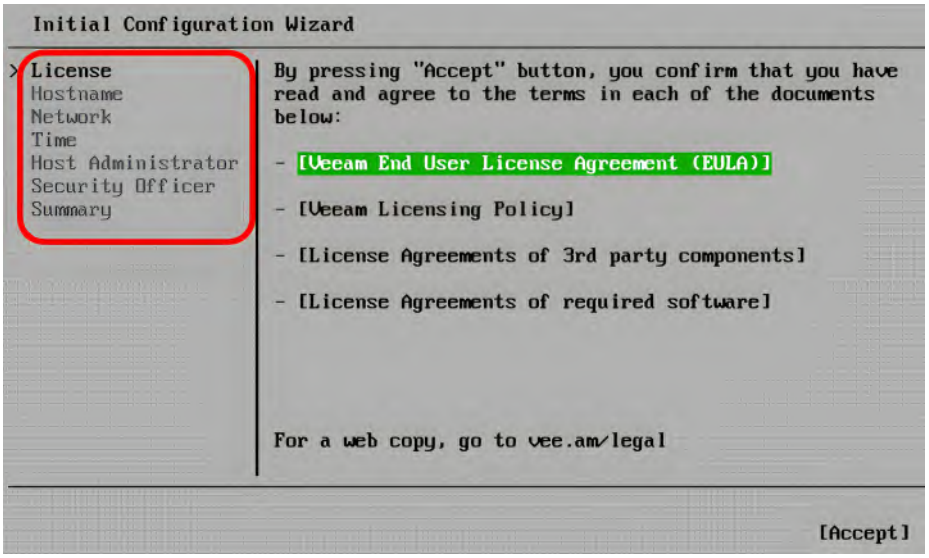


Figure 2.6: Initial Configuration Wizard to complete appliance settings

4. In the **Initial Configuration Wizard**, you complete the following steps:
   Hostname: Type in the Fully Qualified Domain Name (FQDN) of the server, such as `veeam-v13-test.domain.xyz`, as a best practice. (Note that I did not use FQDN during my deployment.)

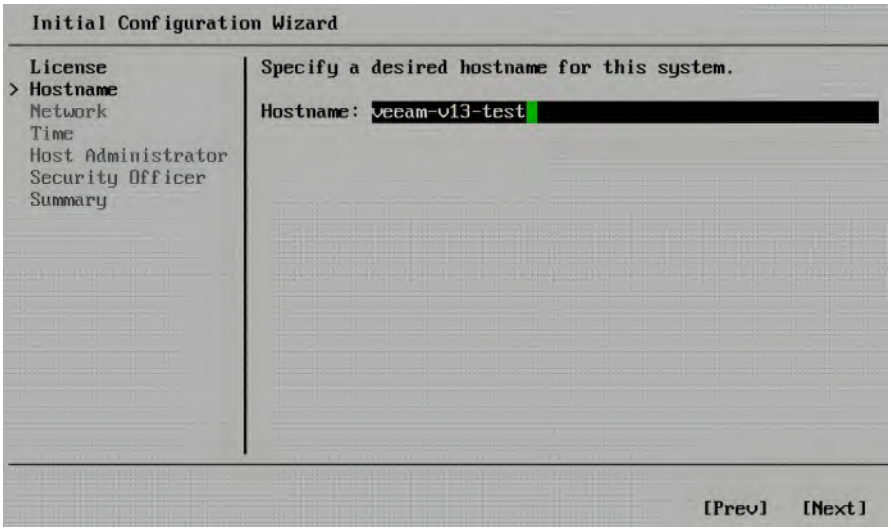


Figure 2.7: Hostname

Network: Set your DHCP/Static IP address here. You can use the **Advanced** options to set other features.

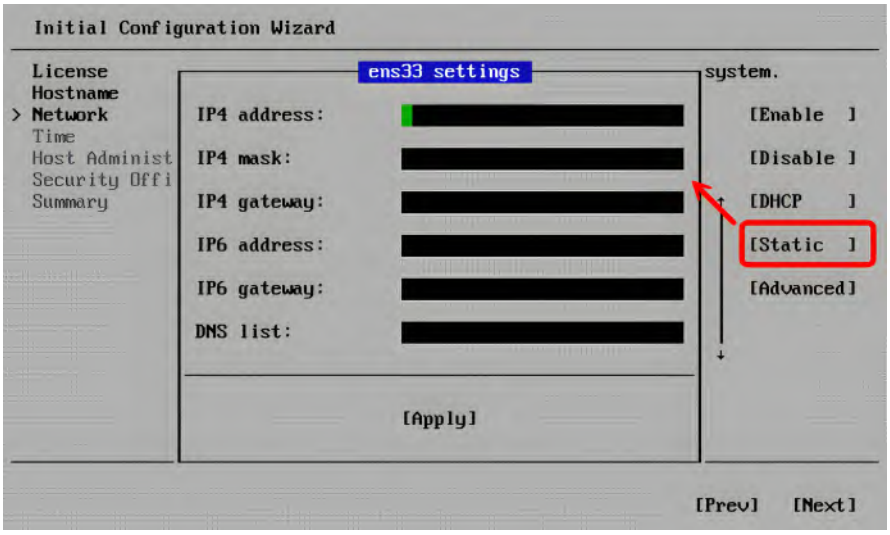


Figure 2.8: Network settings

Time: Set the time servers or server pool and turn on Network Time Security (NTS) if supported.

Technet24

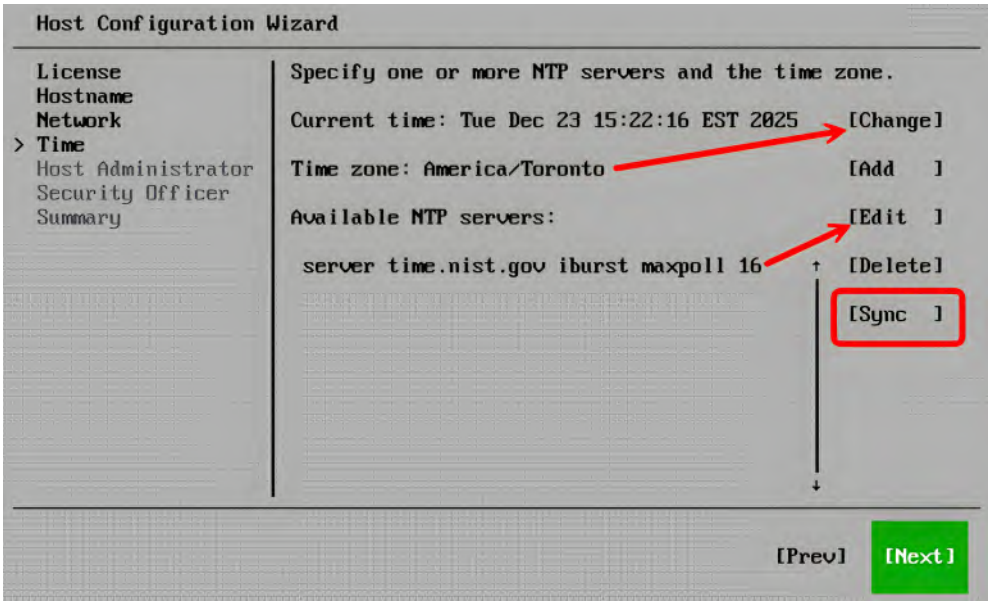


Figure 2.9: Time settings

Host Administrator: The default username is **"veeamadmin"** and set the password.

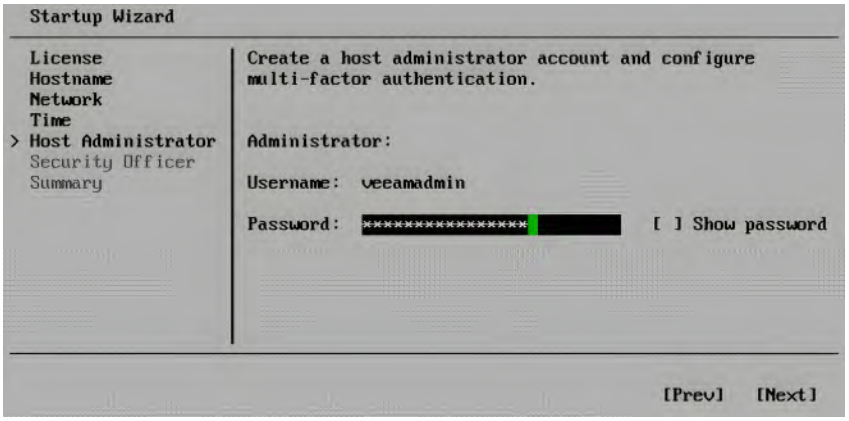


Figure 2.10: Administrator password

Configure the MFA as well.

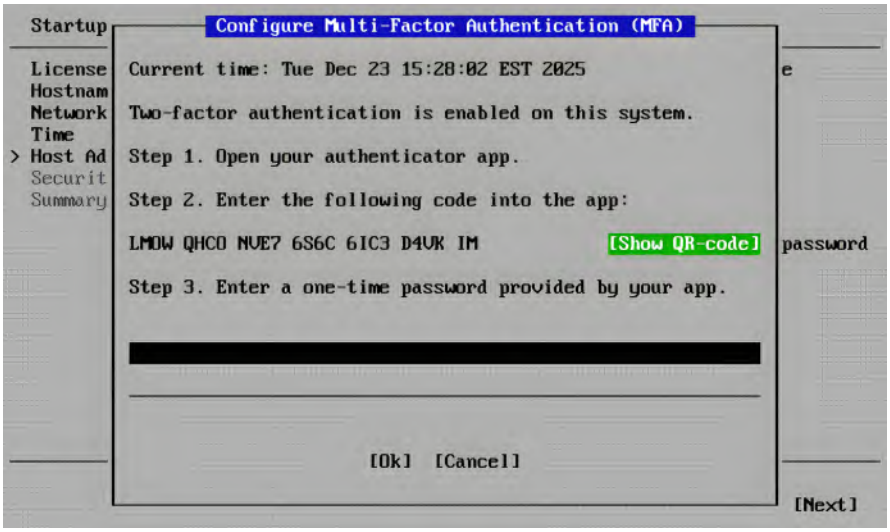


Figure 2.11: MFA setup for Administrator account

Security Officer: Here, you choose to set up the "veeamso" account password, which gets changed on initial login, so make it something easy to remember.

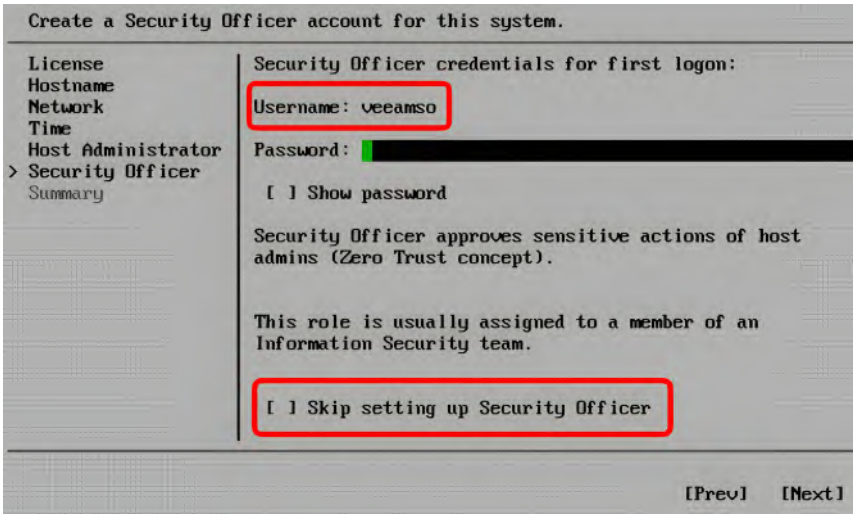


Figure 2.12: Security officer password and option to skip

You can also choose not to set a **Security Officer**; in that case, you are prompted to type `OK` to confirm.

Technet24

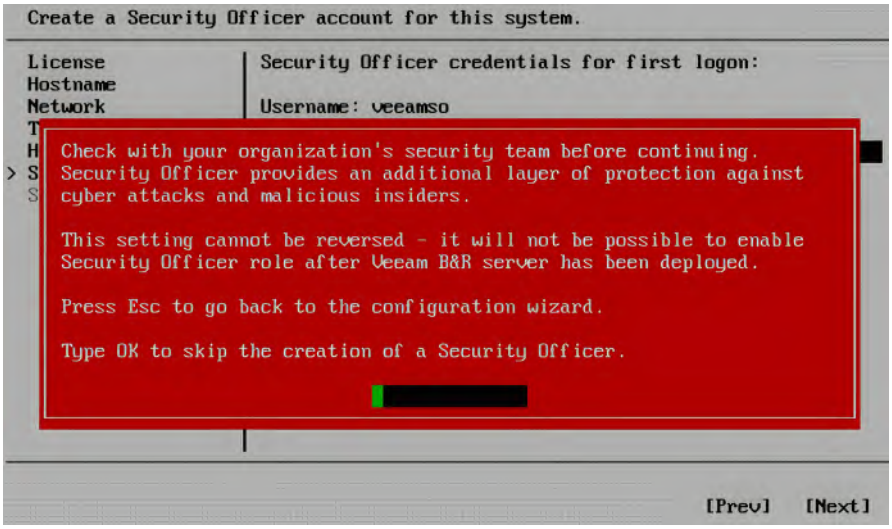


Figure 2.13: Skip security officer setup confirmation

Note

If you plan to use the High Availability (HA) feature mentioned later on in this chapter, you ***MUST*** configure the Veeam Security Officer (veeamso) account at this step; otherwise, it will not be possible to enable HA.

When you set up the **Security Officer** in the **Initial Configuration Wizard**, you must specify a password. On first login to the Security Officer account (veeamso), you are prompted to change the password to a secure value and set up MFA.

Summary: This screen shows you the details of how the server will be configured. Click **Finish** to complete the setup.

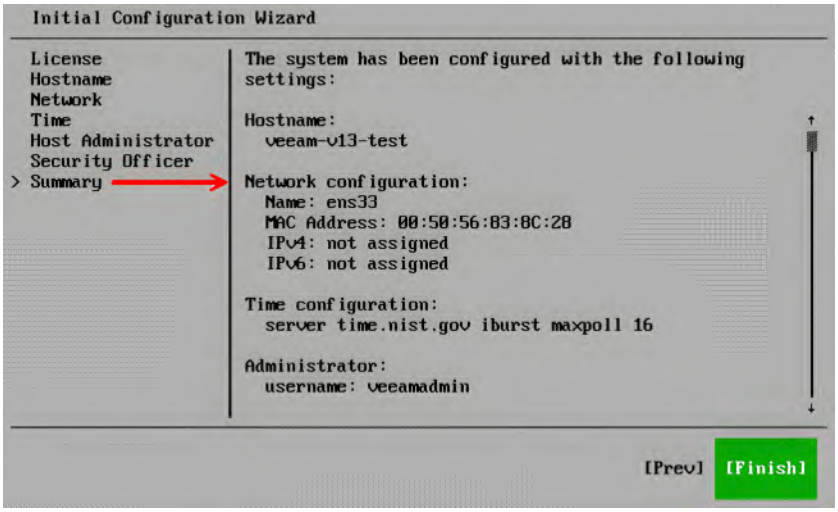


Figure 2.14: Summary screen before finalizing

1. After you complete the **Initial Configuration Wizard**, the system is ready for login and the required information to access the Veeam Software Appliance is displayed

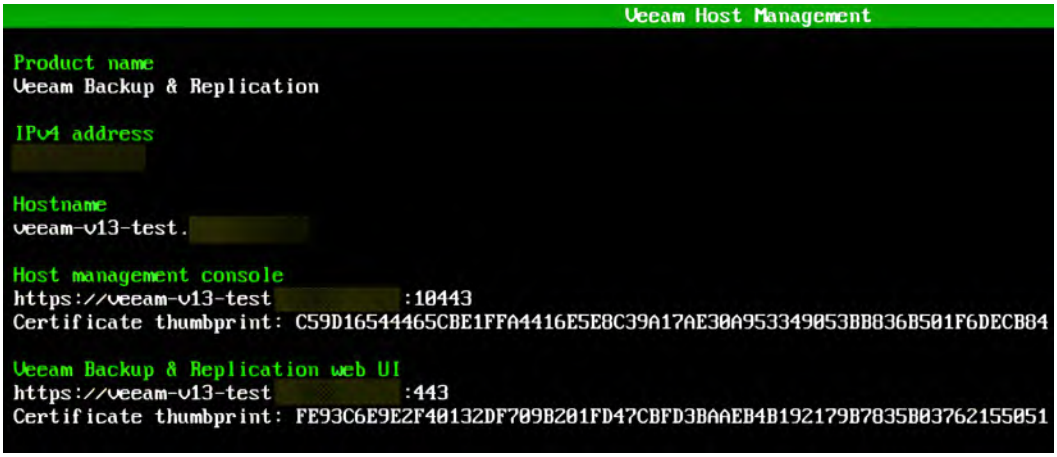


Figure 2.15: Veeam Host Management detail screen

The new Security Officer role provides the following capabilities:

- Approves changes but does not initiate them. You can check out the list of changes the veeamso can make here - `https://helpcenter.veeam.com/docs/vbr/em/hmc_perform_so_tasks.html?ver=13`

Technet24

- Can only be disabled during installation via the Initial Configuration Wizard (not recommended, as it follows security best practices and ensures procedures are followed)
- Protects the system against unauthorized changes

Authorization from the Security Officer is required for the following:

- Root access and password resets
- Restoring configuration from Configuration Backup files
- Advanced deployment using High Availability (HA)
- System monitoring requests, like adding the server to Veeam ONE

There is a migration path from Veeam Backup and Replication 13 on Windows to move to VSA, which involves using the Configuration Backup and a required code from Support.

You must first upgrade to Veeam Backup and Replication v13 on Windows, then migrate to the Veeam Software Appliance. As of now, migration is being performed with Support and will be fully available in version 13.1, scheduled for later in 2026 Q2.

The following are some of the benefits of moving from Windows to VSA:

- Rocky Linux-based appliance – no more Windows licensing costs
- Password length and complexity requirements for better security
- Pre-hardened by default – built according to DISA STIG standards
- MFA enabled for all accounts unless otherwise specified when creating users
- Least privilege permissions service architecture – no SSH or SUDO permissions required
- Updates – managed directly by Veeam and configured in the appliance or via the Windows Console application, including auto updates!
- Inactivity logout for both the web console and command line, which will terminate after 10 minutes

Now that we have reviewed the latest installations of Veeam Backup and Replication v13 software appliance, let's examine how the appliance is security-hardened out of the box and how it uses High Availability for disaster recovery.

Security hardening and high availability

VSA v13 represents a fundamental shift in how organizations deploy and manage their backup infrastructure. Built from the ground up with security as a core principle, the VSA delivers a

pre-hardened, Linux-based platform that eliminates many of the configuration challenges and security risks associated with traditional Windows-based deployments.

Security hardened out of the box

One of the most compelling aspects of VSA v13 is that it arrives pre-hardened and ready for production deployment. Unlike traditional installations where administrators must manually configure and secure the operating system, the VSA implements security best practices automatically from deployment.

Just enough operating system (JeOS) architecture

The foundation of VSA v13's security posture is Veeam JeOS, a purpose-built, minimalistic Linux distribution developed and maintained by Veeam. Based on Rocky Linux, JeOS is specifically designed to streamline the deployment and management of Veeam infrastructure components. The "just enough OS" philosophy includes only the essential services and components required to run Veeam roles, dramatically reducing the attack surface compared with full operating system installations.

Key characteristics of JeOS architecture include the following:

- **Minimal footprint**: Only essential packages and services are installed, eliminating unnecessary software that could introduce vulnerabilities
- **Faster deployments**: The streamlined OS results in quicker installation and boot times
- **Reduced maintenance**: Fewer components mean fewer patches and updates to manage
- **Predictable behavior**: A consistent, known configuration across all deployments

DISA STIG compliance

VSA v13 is pre-hardened according to the **Defense Information Systems Agency** (**DISA**) **Security Technical Implementation Guide** (**STIG**) standards. DISA STIGs are among the most rigorous security configuration guidelines available, initially developed for U.S. Department of Defense systems. By adhering to these standards out of the box, organizations gain enterprise-grade security compliance without the manual effort typically required to configure them.

DISA STIG compliance includes password complexity requirements that must be met during initial configuration. Passwords that do not meet STIG requirements will be rejected during setup, ensuring the appliance starts with strong authentication credentials.

Secure remote access

By default, SSH is disabled on the VSA, eliminating one of the most commonly targeted attack vectors. This design decision reflects best security practices, in which unnecessary services like

Technet24

SSH and root access are disabled unless explicitly required. If SSH access is needed for troubleshooting, it can be temporarily enabled via the web-based management console. Still, the expectation is that normal administrative operations occur through secure web interfaces.

The appliance provides two web-based management interfaces:

- **Host Management Console (Port 10443)**: A cockpit-based interface for system-level administration, including network configuration, user management, and appliance settings

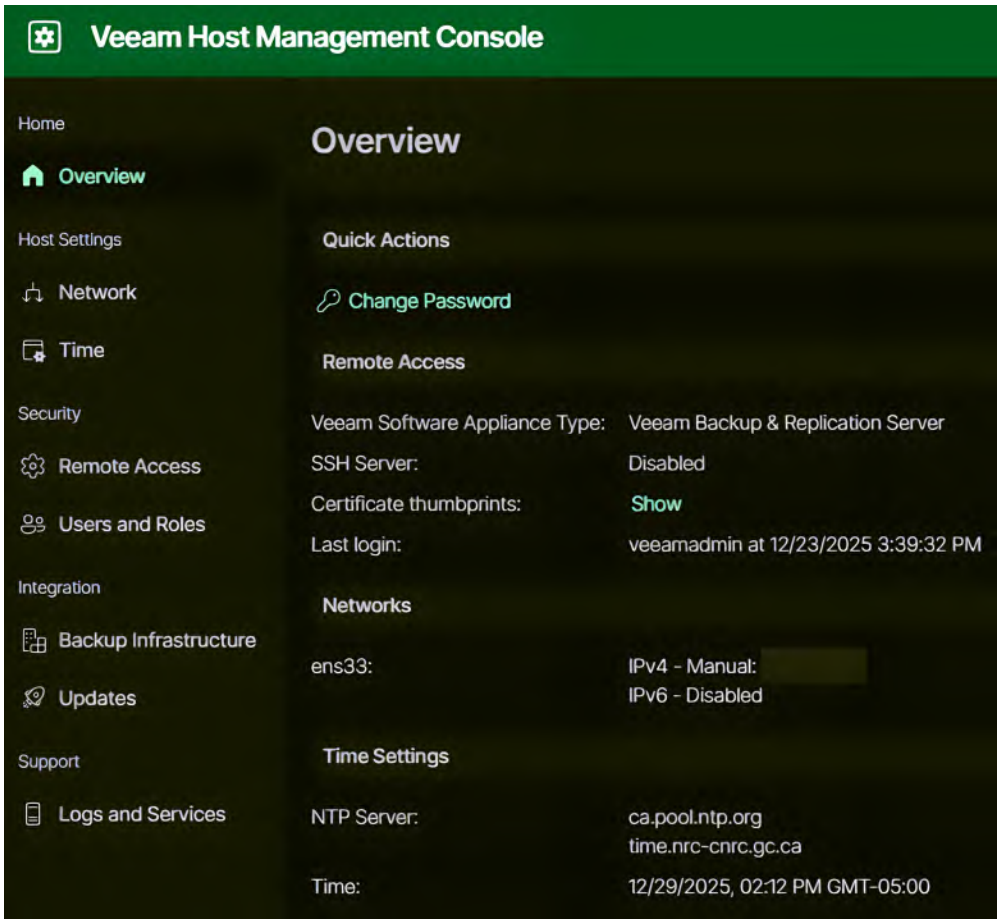


Figure 2.16: Veeam Management Console on port 10443

- **VBR Web UI (Port 443)**: The primary interface for backup and recovery operations, meeting Web Content Accessibility Guidelines (WCAG) standards

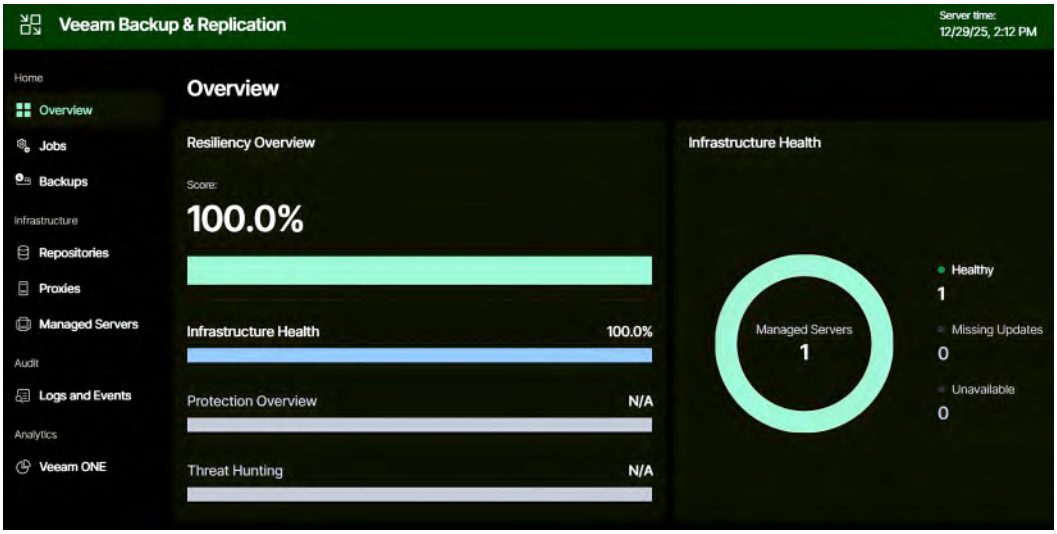


Figure 2.17: Veeam Web UI on port 443

Multi-factor authentication and security officer role

VSA v13 mandates multi-factor authentication (MFA) for administrative access. During initial configuration, administrators must enable MFA using a standard One-Time Password (OTP) app, such as Google Authenticator or Microsoft Authenticator. This requirement applies to both the veeamadmin account and any additional administrator accounts, ensuring that compromised credentials alone are insufficient to gain administrative access.

A significant security enhancement in v13 is the introduction of the Security Officer role. This role provides separation of duties for high-risk operations, requiring approval before actions such as deleting backup data can be executed. The Security Officer account is configured during initial setup. It acts as a check on potentially destructive administrative actions, adding an additional layer of protection against both malicious actors and accidental data loss.

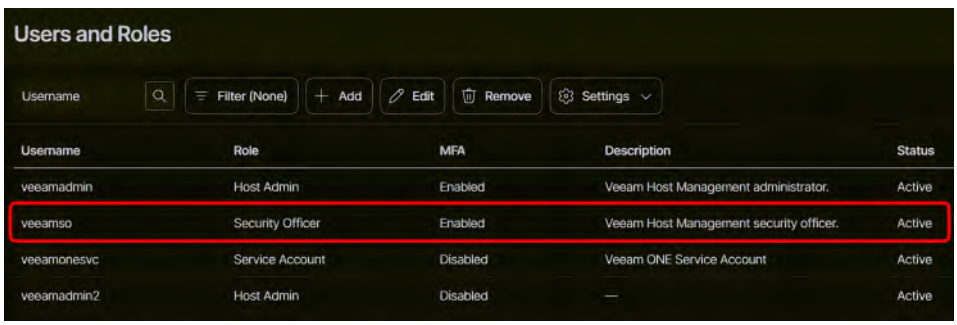


Figure 2.18: Veeam Security Officer account

Technet24

Zero trust architecture

VSA embraces Zero Trust principles throughout its design. This includes RBAC with support for custom roles and scoped permissions, allowing organizations to implement **least-privilege** access policies. Integration with external Identity Providers via SAML 2.0 enables Single Sign-On (SSO) with providers such as Microsoft Entra ID and Okta, centralizing authentication and reducing password sprawl.

Certificate-based authentication is used for component communication and data transit between Veeam components, with data encrypted over TLS. When connecting monitoring tools like Veeam ONE to the VSA, certificate-based pairing can be used instead of storing credentials, further reducing the attack surface.

Automated patching and updates

Veeam manages both operating system updates and appliance patches for the VSA, helping organizations stay secure and up to date with minimal administrative effort. This automatic update approach ensures that security patches are applied consistently and promptly, eliminating the typical scenario in which backup infrastructure lags on critical updates due to management complexities or resource constraints.

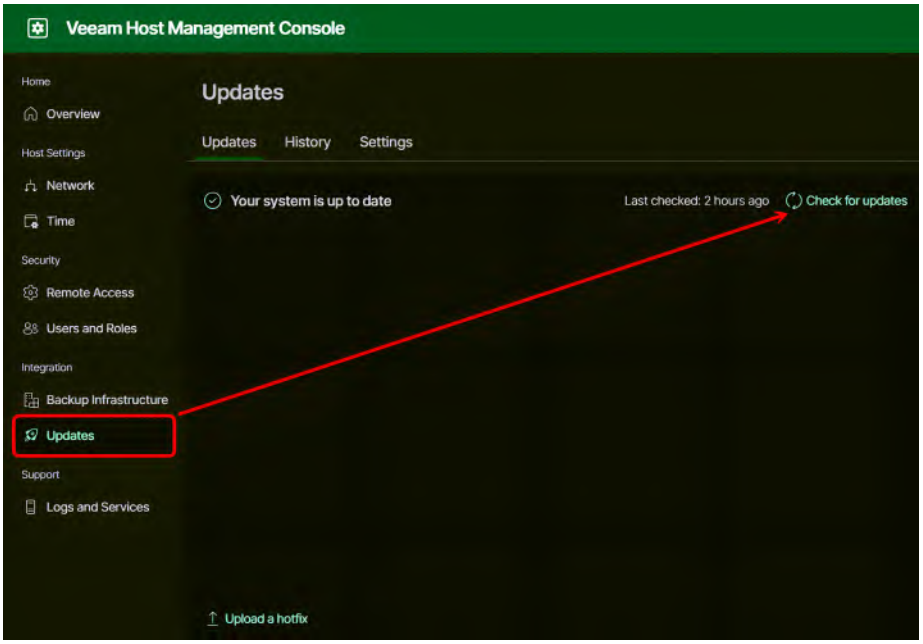


Figure 2.19: Updates screen in the Host Management Console

Figure 2.20 shows the update process delivery from Veeam to the Veeam Software Appliance and all other Veeam Software Appliance devices.

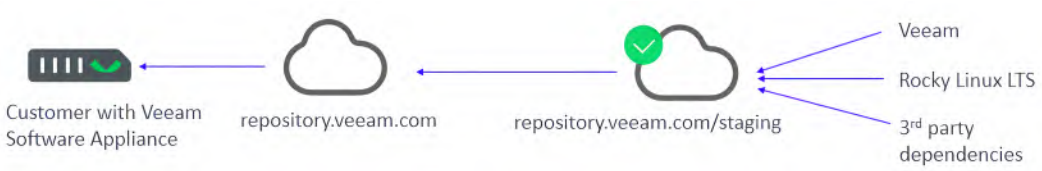


Figure 2.20: Veeam software update process for Veeam Software Appliances

Security features summary

The following table is a summary of security enhancement features that have been added to Veeam Backup and Replication v13.

Security Feature	Implementation
Operating System	Veeam JeOS (Rocky Linux-based minimal OS)
Compliance Standard	DISA STIG compliant out of the box
Access Control	SELinux strict mode, RBAC with custom roles
Authentication	Mandatory MFA, SAML 2.0 SSO support
Remote Access	SSH disabled by default, web-based management
Separation of Duties	Security Officer role for high-risk operations
Updates	Veeam-managed automatic OS and application patches

Table 2.1: Security feature summary

As you can see Veeam has done a great job to enhance security within the new v13 platform. Now let's take a look at the High Availability option that is available for disaster recovery.

High availability

For organizations with strict Recovery Time Objectives (RTOs) and Recovery Point Objectives (RPOs), the backup server is a mission-critical component. VSA v13 introduces native HA clustering capabilities, enabling an active/passive setup that maintains operational continuity during both planned maintenance and unexpected outages.

Technet24

The HA feature in VSA provides redundancy for the backup server component by maintaining two synchronized VSA nodes in an active/passive configuration. The primary node handles all active operations while the secondary node remains synchronized and ready to assume operations if needed.

Cluster architecture

The HA cluster consists of three key components:

- **Primary Node**: The active VSA that handles all backup server operations in read-write mode
- **Secondary Node**: A synchronized standby VSA that maintains a copy of all configuration data in read-only mode
- **Cluster IP**: A virtual IP address that administrators use to connect to the cluster, automatically directing to the current primary node

The synchronization mechanism replicates the entire configuration database across nodes, including credentials, job configurations, and secrets. Importantly, this synchronization can operate even under relatively high network latency, making it suitable for geographically distributed deployments.

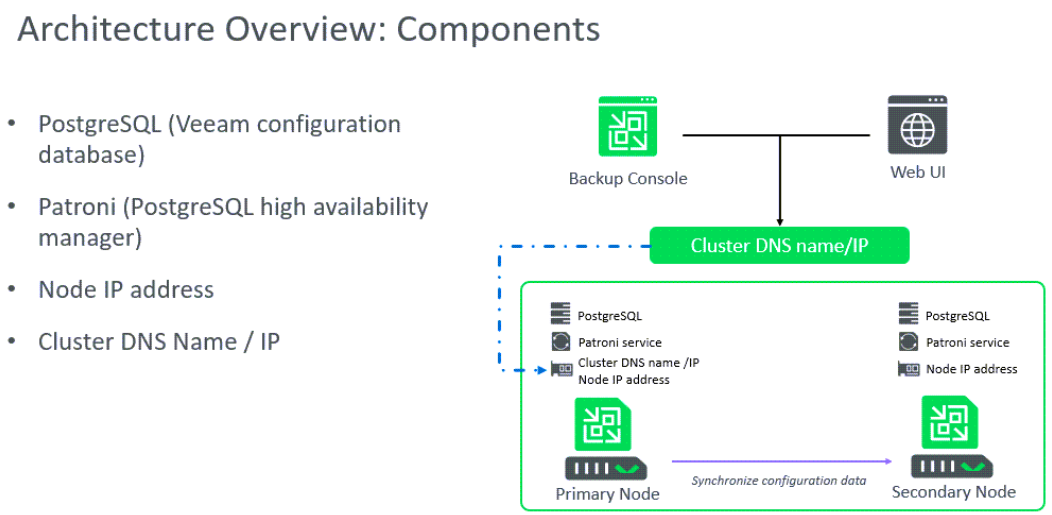


Figure 2.21: VSA HA Architecture Components

To prevent split-brain scenarios in which both nodes might attempt to operate as the primary, Veeam Backup and Replication uses a built-in mechanism: the primary node always operates in read-write mode. In contrast, the secondary node remains in read-only mode. This ensures

that only one node can make configuration changes at any given time, maintaining data consistency across the cluster.

Switchover and failover operations

The HA cluster supports two types of role transitions: switchover for planned maintenance and failover for unexpected outages.

Planned switchover

A switchover is used when both nodes are operational and synchronized, but maintenance is required on the primary node. The process involves the following steps:

1. Ensuring all backup jobs are stopped or completed
2. Initiating the switchover from the VBR console

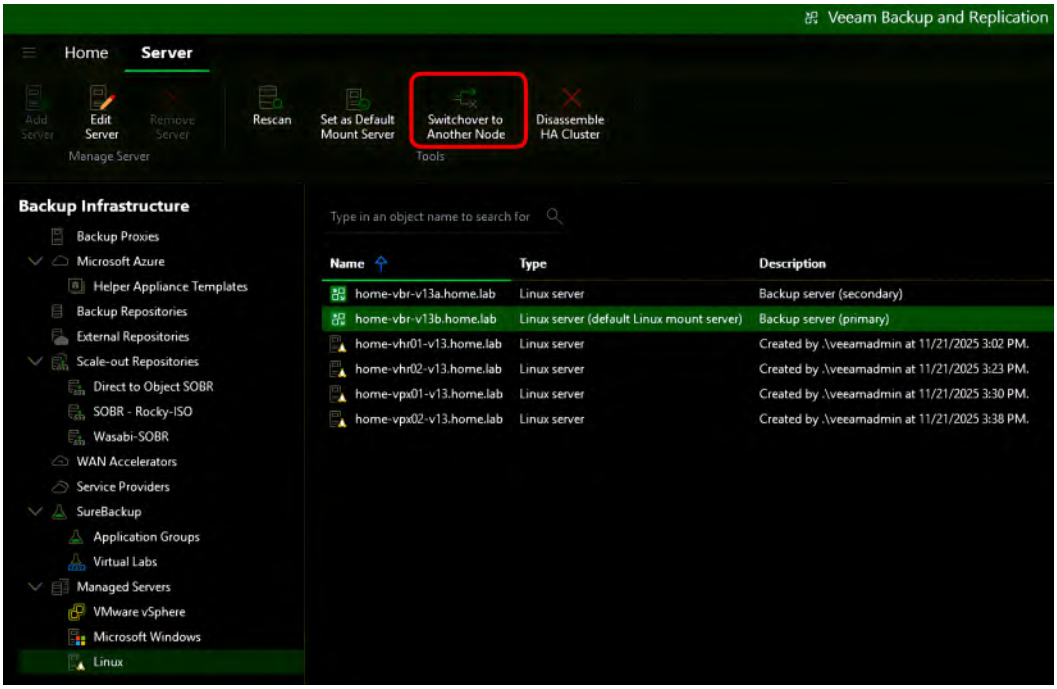


Figure 2.22: Planned switchover option using Veeam Console

3. Waiting for the role transition (typically 10-15 minutes)
4. Reconnecting to the cluster IP to resume operations

Emergency failover

A failover is necessary when the primary node becomes unavailable due to an unexpected outage. Typically, an emergency failover occurs when communication with the cluster's

Technet24

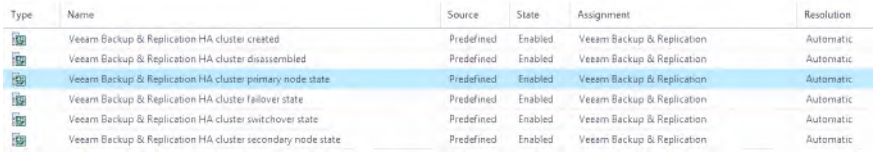
primary node fails. This could be due to networking issues, hardware problems like a VMware host crash or physical host crash, or problems within a virtual machine such as VMware tools, and many more. The process involves the following steps:

1. Attempting to connect to the cluster IP, which prompts that the primary is unresponsive
2. Accepting the option to connect to the secondary node
3. Initiating the failover process to promote the secondary node
4. Reconnecting after promotion completes (approximately 6–10 minutes.)

A critical advantage of the HA architecture is that after failover, there is no need to import configuration backups, recreate jobs, import backups, or re-enter credentials. The backup server continues to operate seamlessly because all configuration data is already synchronized to the secondary node.

Integration with veeam ONE

Veeam ONE v13 provides enhanced monitoring and automation capabilities for HA clusters. The integration includes continuous health monitoring with alarms that notify administrators of health degradation or failover conditions. Additionally, automated remediation actions can trigger a failover when Veeam ONE detects an unhealthy primary node, enabling policy-driven failovers for both planned and unplanned events.



Type	Name	Source	State	Assignment	Resolution
	Veeam Backup & Replication HA cluster created	Predefined	Enabled	Veeam Backup & Replication	Automatic
	Veeam Backup & Replication HA cluster disassembled	Predefined	Enabled	Veeam Backup & Replication	Automatic
	Veeam Backup & Replication HA cluster primary node state	Predefined	Enabled	Veeam Backup & Replication	Automatic
	Veeam Backup & Replication HA cluster failover state	Predefined	Enabled	Veeam Backup & Replication	Automatic
	Veeam Backup & Replication HA cluster switchover state	Predefined	Enabled	Veeam Backup & Replication	Automatic
	Veeam Backup & Replication HA cluster secondary node state	Predefined	Enabled	Veeam Backup & Replication	Automatic

Figure 2.23: Veeam ONE built-in HA Cluster alarms

Note

When adding an HA Cluster to Veeam ONE, the Analytics service will be installed on each node if not currently present. You need to ensure that the same user credentials are created on each node with the same password so that operational monitoring by Veeam ONE will work without issue.

High Availability functionality has specific licensing and version requirements:

- **VSA Version**: Veeam Software Appliance 13.0.1 or later
- **Licensing**: Veeam Data Platform Premium Edition (includes VBR, Veeam ONE, and Veeam Recovery Orchestrator.)

- **Infrastructure**: Two VSA nodes with network connectivity
- **Network**: Reserved cluster IP address with proper DNS resolution – both nodes must be in the same Subnet/L2 Network as per `https://helpcenter.veeam.com/docs/vbr/userguide/high_availability_limitations.html?ver=13#ha-cluster-network-limitations`
- **Security Officer**: If configured, Security Officer approval is required to enable HA

Note

HA is the only major v13 feature that requires **VDP Premium Edition**. All other new features are available in the Foundation Edition.

Benefits of high availability

Implementing HA for the VSA provides numerous operational and business benefits:

- **Reduced downtime**: Minimizes backup server unavailability during both planned maintenance and unexpected failures
- **Fast recovery**: Eliminates the need to restore configuration backups or recreate jobs after a failure
- **Simplified DR**: Removes the requirement for a separate standby backup server with manual configuration restoration
- **Automatic updates**: Both nodes receive updates automatically, maintaining consistency across the cluster
- **Geographic distribution**: Nodes can synchronize across high-latency links, enabling cross-site redundancy
- **Operational continuity**: Backup and recovery operations can continue without significant interruption
- **Credential preservation**: All credentials and secrets are synchronized, eliminating re-authentication requirements
- **Proactive monitoring**: Integration with Veeam ONE enables automated health monitoring and policy-driven failovers

Note

Important considerations:

While HA provides excellent protection for the backup server component, it is crucial to understand its scope. The HA feature synchronizes configuration and operational data

Technet24

> between nodes; it is not a replacement for backup. If issues such as corruption or misconfiguration occur on the primary node, they will also be replicated to the secondary node.
>
> For repository-level high availability, Veeam recommends using S3-compatible object storage with node-loss protection or configuring Scale-Out Backup Repositories (SOBR) to continue backups if an extent is unavailable. These approaches complement VSA HA by providing storage-tier resilience.

Now that we have covered the new security hardening and high-availability clustering options, let's look at additional features in the web-based management interface and new enhanced RBAC roles.

Web-based management interface and enhanced RBAC

VSA v13 introduces a modern, web-based user interface alongside significantly enhanced Role-Based Access Control (RBAC) capabilities. These features represent a fundamental shift in how administrators interact with and secure their Veeam backup infrastructure, eliminating traditional dependencies on Windows-based thick clients while providing granular control over user permissions.

The new web-based user interface

One of the most significant additions in VSA v13 is the introduction of a comprehensive web-based user interface. This modern interface eliminates the need for legacy consoles or thick clients, empowering administrators to securely manage data protection workflows from anywhere with a web browser. The Web UI complies with the Web Content Accessibility Guidelines (WCAG), ensuring accessibility for diverse user needs.

Key features and capabilities

The Web UI provides administrators with a rich set of features designed for modern backup management:

- **Built-in monitoring dashboard**: Real-time visibility into critical metrics and backup infrastructure health, enabling rapid alert resolution and proactive management
- **Veeam ONE integration**: Seamless integration with Veeam ONE directly within the Web UI, providing enhanced observability and reporting capabilities without switching between applications

- **Dark mode support**: A refreshed UI design with a dark mode option, reducing eye strain during extended management sessions
- **SAML 2.0 Single Sign-On**: Integration with enterprise identity providers such as Microsoft Entra ID and Okta for centralized user authentication
- **Console download**: The backup console setup can be downloaded directly from the backup server using a link at the bottom of the Web UI logon page
- **Responsive design**: Clean, responsive interface supporting management from data centers, remote offices, or mobile devices

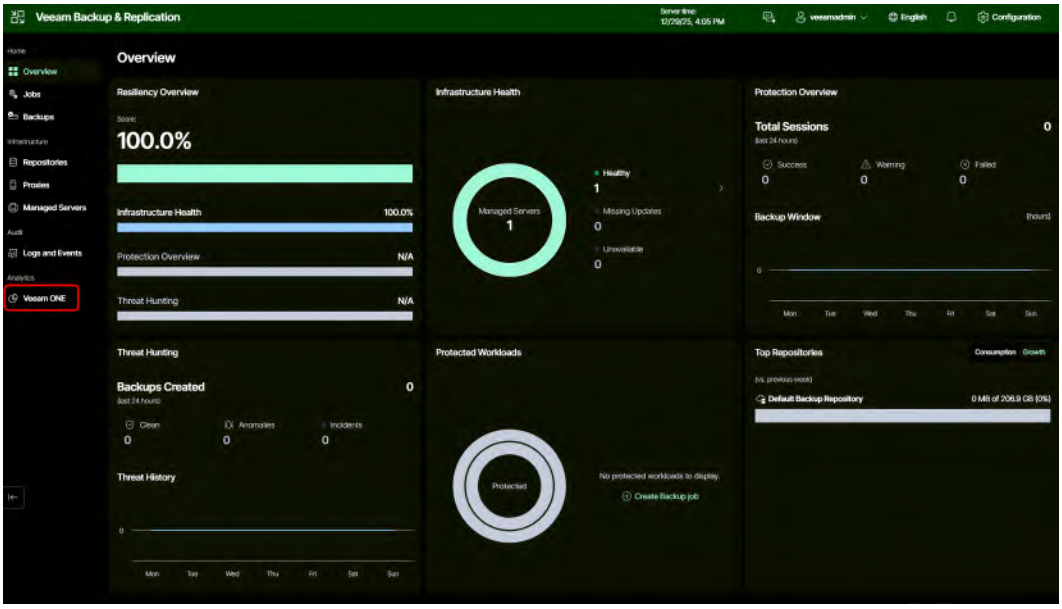


Figure 2.24: Redesigned Web-Based management interface

Web-based approach brings several benefits:

- **Platform independence**: Manage backup infrastructure from any operating system with a modern web browser, eliminating Windows workstation requirements for daily administration
- **Reduced attack surface**: No need to install thick clients on administrator workstations, minimizing potential security vulnerabilities
- **Simplified access**: Administrators can access the management interface from anywhere without VPN dependencies when properly secured

Technet24

- **Consistent experience**: Unified interface across deployments ensures consistent operational procedures and reduced training requirements
- **Zero Trust alignment**: Combined with SAML SSO and RBAC, the Web UI supports modern zero trust security architectures

Note

In the initial v13 release, some advanced configuration tasks may still require the traditional Veeam Console. Feature parity between the Web UI and thick client continues to expand with each update.

Enhanced Role-Based access control (RBAC)

VSA v13 delivers a completely redesigned RBAC framework that enables organizations to implement granular, scoped roles that define exactly what operators can view, back up, or restore. This enhanced model supports the principle of least privilege and aligns with increasingly stringent security and zero-trust regulations.

Built-in roles

VSA v13 includes several predefined roles that cover common operational scenarios as shown in *Table 2.2*.

Role	Description and Permissions
Veeam Backup Administrator	Full access to all Veeam Backup and Replication functionality, including backup creation, data recovery, job management, user role management, and system configuration.
Veeam Backup Operator	Create and manage backup jobs, modify backup settings, start and stop jobs, and view backup history. Cannot recover data or manage user roles.
Veeam Restore Operator	Recover data from backups, including browsing backups, selecting recovery points, and restoring files or objects. Cannot create backups or manage jobs.
Veeam Backup Viewer	Read-only access to view the console, backup jobs, backups, and recovery information. Cannot create backups, manage jobs, or recover data.

Role	Description and Permissions
Security Administrator	Manage credentials (add, edit, delete), configure the Security and Compliance Analyzer, and approve Four-Eyes Authorization requests—granting console access with viewing permissions for most objects.
Security Officer	Host Management UI role (Linux/VSA only) that approves sensitive management functions such as High Availability configuration. Provides protection against malicious system administration. MFA is always enforced.
Veeam Tape Operator	Manage tape devices, create tape backup jobs, backup and restore from tapes, and view tape inventory. Cannot create disk backups or manage user roles.
Incident API Operator	No console access. Interacts only with the Veeam Incident API REST endpoint to manage malware detection events. Functions as a service account with MFA disabled by default.

Table 2.2: Veeam built-in roles

Custom role creation

One of the most powerful enhancements in v13 is the ability to create custom roles with precisely scoped permissions. The new Custom Role Wizard provides an intuitive interface for defining roles and their associated permissions, enabling organizations to implement the principle of least privilege at a granular level.

Technet24

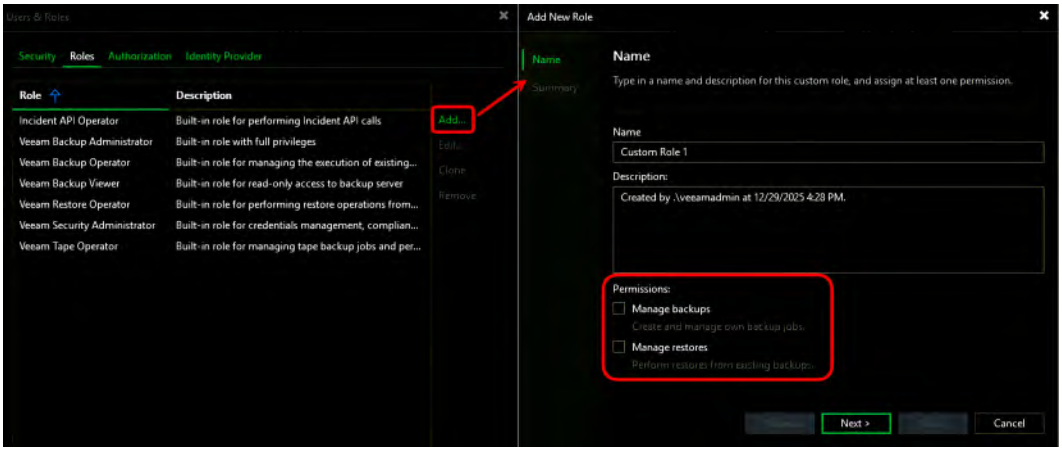


Figure 2.25: Adding a custom role

The following shows some scenario examples of typical roles that would be configured in a Veeam environment, with a use case, configuration, and the resulting access.

Examples of role configurations

The following examples demonstrate practical applications of custom RBAC configurations.

Scenario 1: Departmental backup operator

Use Case: A large organization wants to delegate backup management for the Finance department to a dedicated team member without exposing other workloads.

Configuration:

- **Role Type**: Backup Operator
- **Backup Scope**: Finance vCenter folder only
- **Restore Scope**: None (separate restore role assigned to a different team)

Result: User can create, modify, and run backup jobs for Finance VMs only. They cannot view or manage backups for other departments, nor perform restore operations.

Scenario 2: help desk restore specialist

Use Case: The IT help desk needs to perform file-level restores for end users but they should not be able to restore entire VMs or access production infrastructure.

Configuration:

- **Role Type**: Restore Operator
- **Restore Scope**: All workstation backups

- **Target Restrictions**: Original location only
- **Permitted Operations**: Guest OS file restores only

Result: Help desk staff can browse workstation backups and restore individual files to their original locations. They cannot perform VM-level restores, restore them to alternate locations, or access server backups.

Scenario 3: DR recovery coordinator

Use Case: A disaster recovery team member needs full restore capabilities for critical workloads to the designated DR infrastructure, without modifying backup jobs or configurations.

Configuration:

- **Role Type**: Restore Operator
- **Restore Scope**: Tier 1 critical applications (defined by VMware tags)
- **Target Restrictions**: DR site vCenter cluster only
- **Permitted Operations**: Full VM restore, Instant VM Recovery

Result: DR coordinator can perform full VM restores and Instant Recovery for Tier 1 applications exclusively to the DR site. They cannot modify backup policies or restore production infrastructure.

Scenario 4: compliance auditor

Use Case: External auditors need to verify backup compliance without the ability to modify any configurations or access backup data.

Configuration:

- **Role Type**: Veeam Backup Viewer (built-in)
- **Scope**: All backup jobs and repositories
- **Access Method**: SAML SSO with time-limited access

Result: Auditors can view job configurations, backup schedules, retention settings, and success/failure metrics. They cannot initiate backups, perform restores, or change any settings.

As you can see with these scenarios, assigning roles can help with both access and security to ensure users do not have too many permissions. Now let us look at the security benefits when using the new enhanced RBAC within Veeam v13.

Technet24

Security benefits of enhanced RBAC

The enhanced RBAC framework in VSA v13 delivers significant security improvements:

- **Principle of Least Privilege**: Users receive only the minimum permissions required for their role, reducing the blast radius of compromised accounts
- **Separation of duties**: Critical operations can be split across multiple roles, requiring coordination for sensitive actions and preventing a single point of failure in security
- **Four-eyes authorization**: High-risk operations can require approval from the Security Administrator before execution, providing an additional verification layer
- **Security Officer approval**: Sensitive management functions on the VSA require Security Officer approval, protecting against malicious administrative actions
- **Mandatory MFA**: Security Officer accounts always require multi-factor authentication, and MFA can be enforced for other privileged roles
- **Compliance support**: Granular role definitions help organizations meet regulatory requirements such as SOC 2, HIPAA, and GDPR that mandate access controls and separation of duties
- **Zero Trust alignment**: Combined with SAML SSO and the hardened VSA platform, the RBAC framework supports modern zero trust architectures where no user or system is implicitly trusted

The introduction of a modern Web UI and enhanced RBAC capabilities in Veeam Software Appliance v13 represents a significant advancement in backup infrastructure management. The Web UI provides platform-independent access with enterprise-grade features. At the same time, the enhanced RBAC framework enables organizations to implement precise access controls that align with security best practices and compliance requirements. Together, these features position VSA v13 as a mature, enterprise-ready backup solution that meets the demands of modern IT environments and regulatory landscapes.

We have now reviewed the new web-based management interface and the new Enhanced RBAC controls. The final topic we will cover in this chapter is the latest auto-updates feature in the VSA.

Centralized update management and auto-updates

One of the most significant operational improvements introduced with the Veeam Software Appliance in version 13 is the centralized update management system. This new architecture fundamentally changes how updates are distributed and applied across your Veeam infrastructure, delivering enhanced security, reduced administrative overhead, and improved compliance posture.

Understanding the update architecture

The VSA 13 update system is managed by the **Veeam Updater** service, which runs on the backup server and all infrastructure components deployed from Veeam appliances. This service communicates with components through the Veeam Backup and Replication REST API and uses the Veeam Identity Service for authorization. By default, updates are downloaded directly from the official Veeam repository at `repository.veeam.com/vsa`, though organizations can configure local mirror repositories for air-gapped environments.

The Veeam Updater service checks for available updates every 12 hours and automatically updates itself to ensure compatibility with the latest update packages. This self-updating capability ensures the update mechanism remains current and can handle new update formats and security requirements.

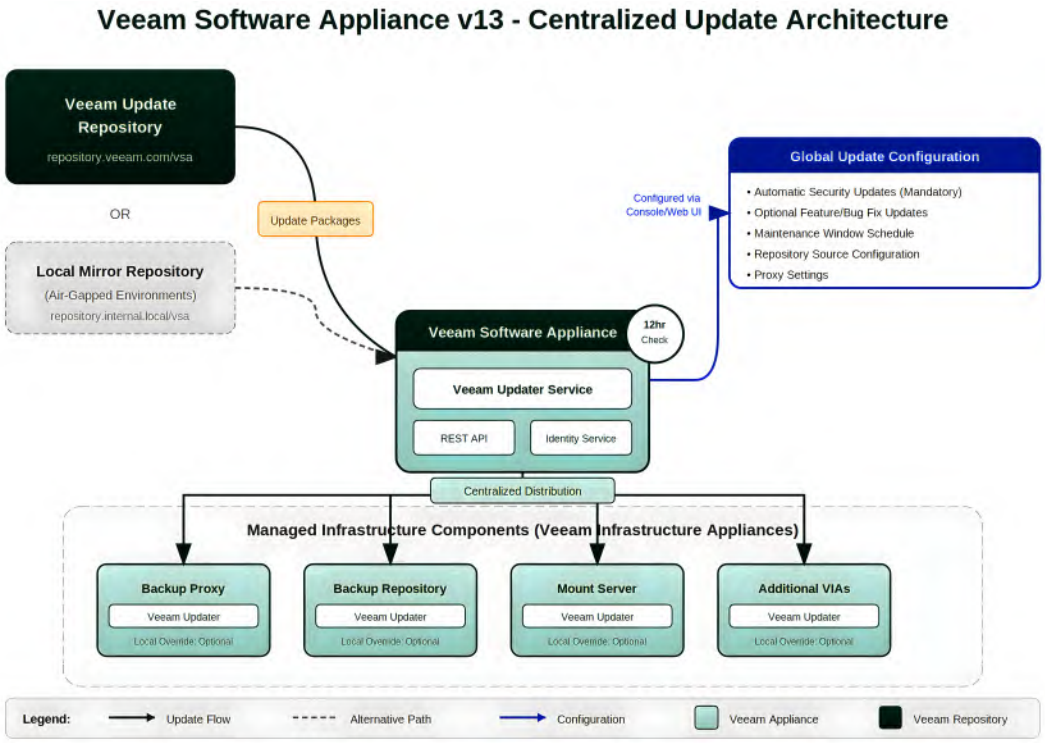


Figure 2.26: Veeam Updater process flow

Technet24

Update categories and policies

VSA 13 categorizes updates into distinct types, each with specific policies governing their installation behavior. Understanding these categories is essential for maintaining a secure and compliant backup infrastructure.

Mandatory updates

Specific updates are classified as mandatory and are installed automatically without user intervention. These cannot be skipped, deselected, or canceled:

- **Operating system security updates**: Critical security patches for the underlying JeOS are applied automatically to address vulnerabilities and maintain system integrity
- **Veeam Backup & Replication security updates**: Security fixes for the backup software itself are installed immediately to protect against known vulnerabilities and potential exploits

This mandatory update policy represents a significant shift from the Windows-based deployment model, where administrators had complete control over patch timing. The rationale is straightforward: in an appliance model where Veeam manages the entire stack, ensuring security updates are applied promptly is both a responsibility and a capability that benefits all customers.

Optional updates

While security updates are mandatory, feature and bug fix updates provide flexibility for organizations that require more controlled rollout schedules:

- Feature updates: New capabilities and enhancements can be enabled for automatic installation or applied manually during planned maintenance windows. An example would be an update to the `.NET` Framework installed on the appliances.
- Bug fix updates: Non-security bug fixes follow the same optional policy, allowing administrators to evaluate changes before deployment. An example would be a fix addressing an issue identified by support, where specific files are patched to resolve the problem
- Minor releases (such as 13.x): Point releases containing cumulative fixes and minor features are optional by default.
- Major releases (such as V14.0): Major version upgrades require explicit administrator action and are never applied automatically.
- Private fixes: Targeted fixes provided by Veeam Support for specific issues are optional and typically applied manually. For example, if an issue is found within a feature like

High Availability, Veeam may send you a package that patches specific files within the appliance to resolve it.

Global update settings

The global update configuration is applied to all backup infrastructure components deployed from the Veeam Software Appliance or the Veeam Infrastructure Appliance. These settings can be configured in either the Veeam Backup and Replication console or the Web UI, and changes synchronize automatically between the two interfaces.

Follow these steps to configure global update settings using the console:

1. Navigate to the main menu and select **Updates | Update Settings**
2. On the **General** tab, configure the update installation options. By default, only security updates are installed automatically, but you can expand this to include optional updates (excluding major releases).

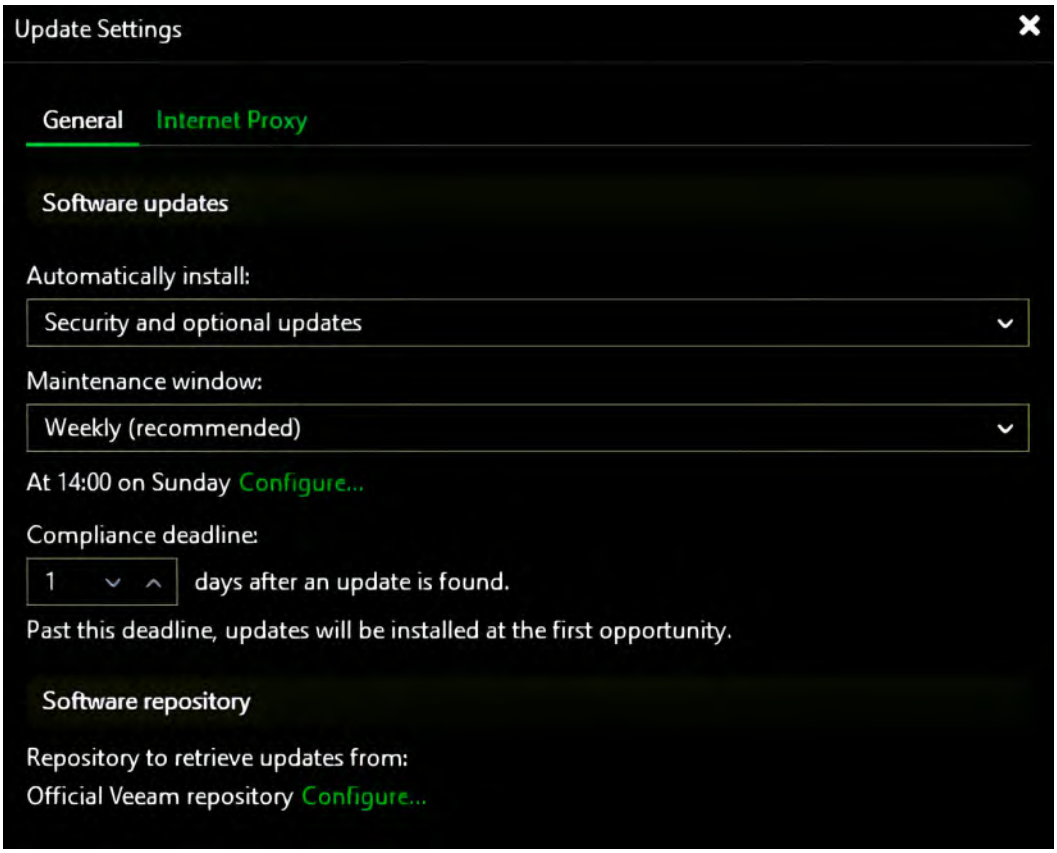


Figure 2.27: Global Update Settings

Technet24

3. Configure the **Maintenance window** to schedule automatic installations on a weekly or monthly basis. If you prefer manual installation, select **None**. The schedule uses the backup server's time zone.

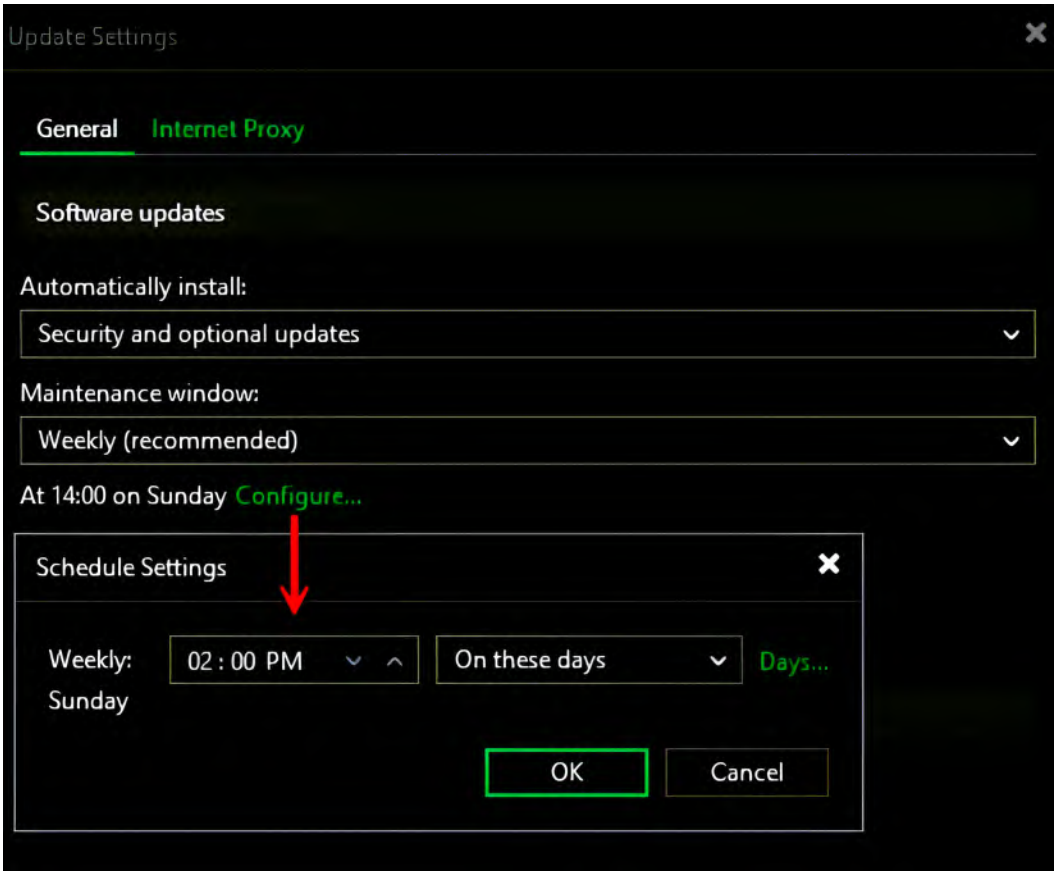


Figure 2.28: Maintenance window and schedule configuration

4. Specify the **Software repository** location. The default is the official Veeam repository, but you can configure a local mirror if required.

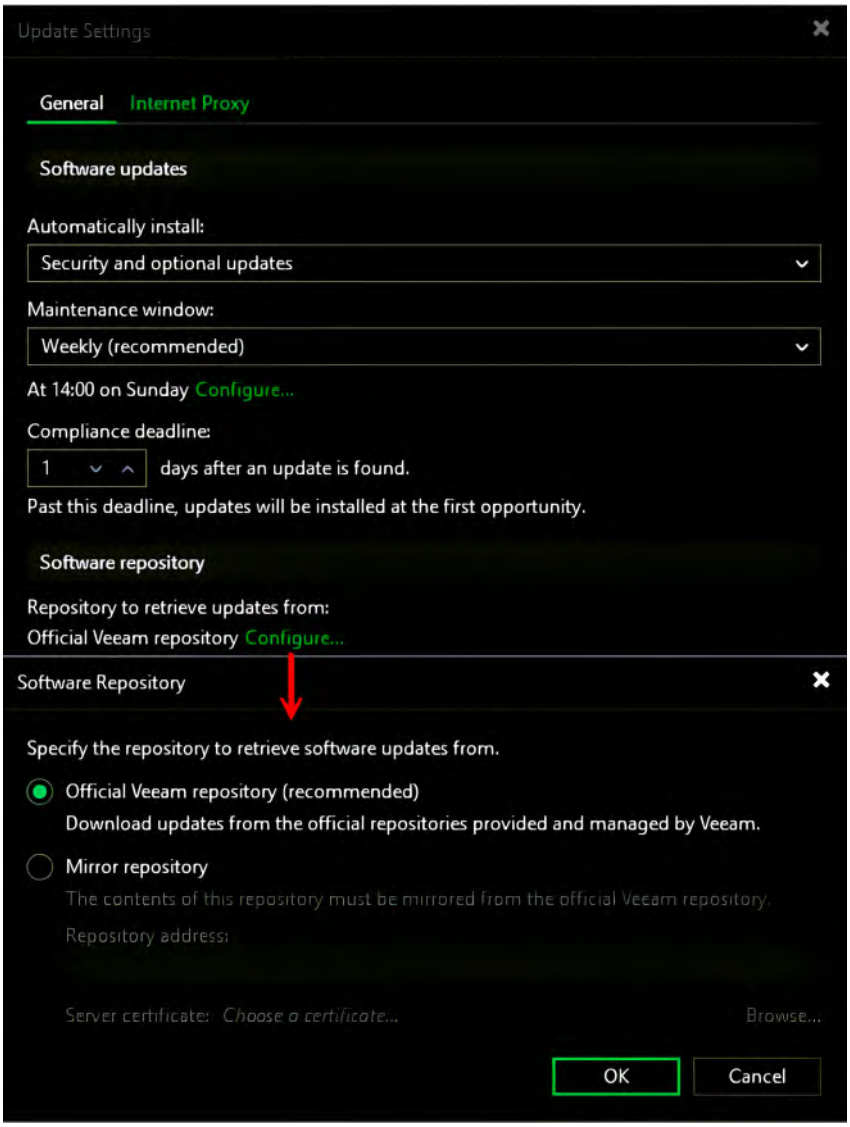


Figure 2.29: Software repository configuration

5. On the **Internet Proxy** tab, add proxy server details if your environment requires one for internet access

These steps outlined are how you configure the auto update policy for all the appliances within your Veeam infrastructure. Now, let us look at how you can override this setting if specific components require a different update configuration.

Technet24

Component-specific configuration

In some scenarios, individual infrastructure components may require different update settings than the global policy. The custom update configuration provides this flexibility, with component-specific settings taking priority over global settings.

Follow these steps to configure custom update settings for a specific component:

1. Log in to the **Veeam Host Management Console** installed on the target infrastructure component
2. Navigate to **Settings** and click **Update Settings**
3. Enable **Manual settings override** to activate component-specific configuration

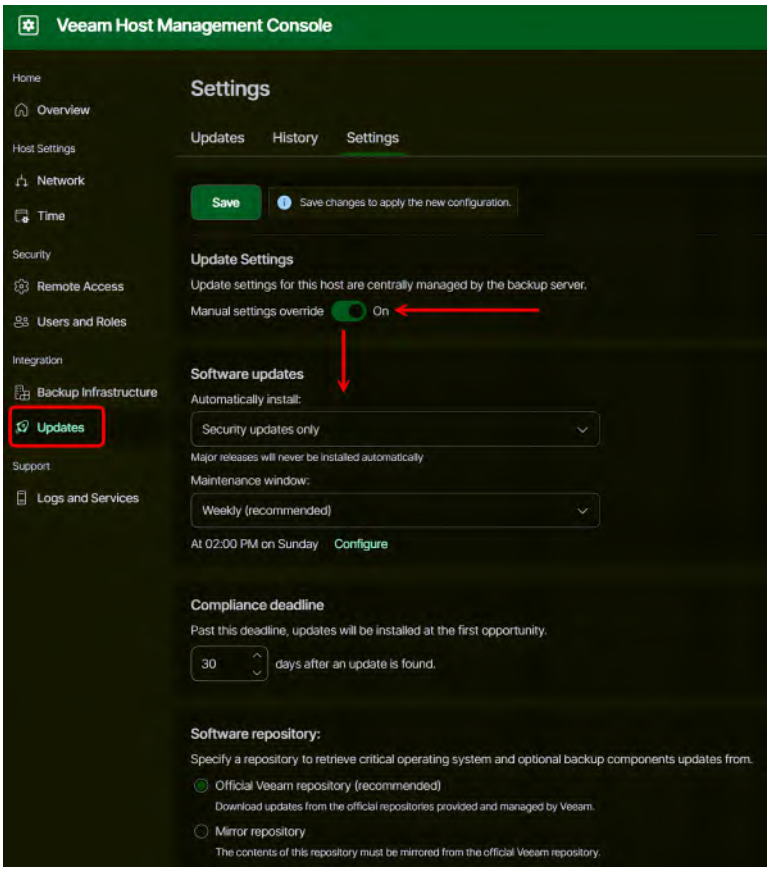


Figure 2.30: Manual override of Update settings on appliance

4. Configure the desired update options, maintenance window, and repository settings specific to this component

This capability is beneficial for components in different time zones, remote locations with limited bandwidth, or environments with specific compliance requirements that mandate separate update schedules.

Configuring updates for air-gapped environments

Organizations operating in air-gapped or restricted network environments can still benefit from centralized updates by deploying a local update repository mirror. This approach ensures all Veeam infrastructure components can receive updates without requiring direct internet access.

Follow these steps to configure a local mirror repository:

1. Deploy a web server within your network that can serve as the update repository mirror
2. Periodically synchronize update packages from the official Veeam repository (`repository.veeam.com/vsa`) to your local mirror using appropriate file transfer methods
3. Access the Veeam Host Management console for your VSA and navigate to **Settings** | **Updates**
4. Enable **Manual settings override** and configure the **Software Repository** to point to your local mirror (e.g., `https://repository.internal.local/vsa`)

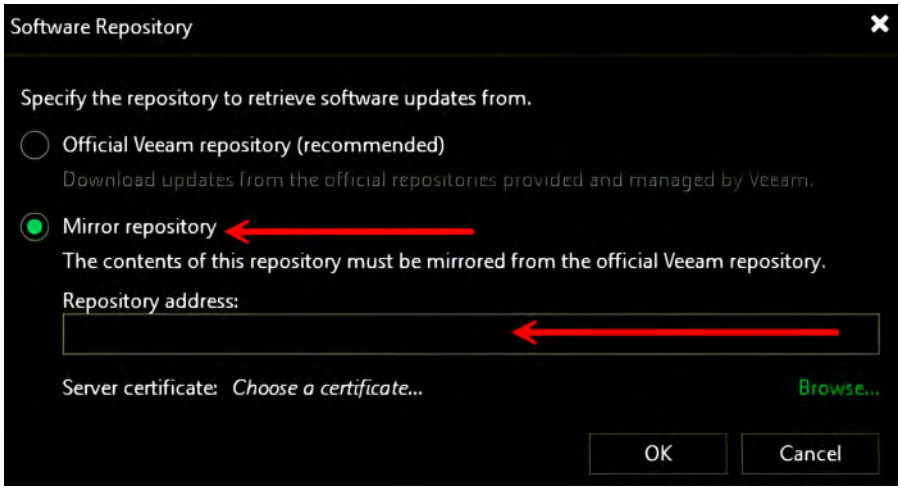


Figure 2.31: Mirror repository URL

5. For HTTPS repositories, ensure the appropriate SSL certificate is configured

Technet24

Update types summary

The following table summarizes the different update categories and their default behavior:

Update Type	Installation Policy	Administrator Control
OS Security Updates	Mandatory	Cannot be disabled or skipped
VBR security updates	Mandatory	Cannot be disabled or skipped
Feature Updates	Optional	Can enable auto-install or apply manually
Bug Fix Updates	Optional	Can enable auto-install or apply manually
Minor Releases (13.x)	Optional	Can enable auto-install or apply manually
Major Releases (V14.0)	Manual only	Never auto-installed; requires explicit action
Private Fixes	Manual only	Applied as needed per Veeam Support guidance

Table 2.3: Update type summary

Benefits of centralized update management

The centralized update architecture in VSA 13 delivers substantial benefits across security, operations, and compliance domains.

- **Enhanced security posture**
    - **Rapid vulnerability remediation**: Security updates for both the operating system and Veeam software are applied automatically, eliminating the delay between patch release and deployment that often exists in manually managed environments

    - **Consistent security baseline**: All appliances in your infrastructure maintain the same security patch level, preventing scenarios where some components remain vulnerable while others are patched
    - **Reduced attack surface**: The automatic patching model significantly reduces the window of exposure for known vulnerabilities, a critical consideration for backup infrastructure that protects business-critical data
- **Reduced administrative overhead**
    - **Unified management interface**: Configure update policies once at the global level and have them automatically propagate to all managed appliances, eliminating the need to manage updates on each component individually
    - **Scheduled maintenance windows**: Define when updates should be applied, allowing installations to occur during low-activity periods without manual intervention
    - **Self-maintaining updater**: The Veeam Updater service updates itself automatically, ensuring the update mechanism remains functional without requiring separate maintenance
- **Improved compliance and auditability**
    - **Demonstrable patch management**: The automatic application of security updates provides auditable evidence that systems are maintained according to security best practices
    - **Update history tracking**: Complete visibility into what updates have been applied and when, supporting compliance reporting and audit requirements
    - **Policy enforcement**: Mandatory security updates ensure that compliance requirements around timely patching are met without relying on manual processes
- **Operational continuity**
    - **Vendor-managed stack**: With Veeam managing both the application and underlying operating system updates, compatibility between layers is guaranteed, eliminating a common source of update-related issues
    - **Controlled feature rollout**: While security updates are mandatory, feature updates remain optional, allowing organizations to control when new functionality is introduced to their environment
    - **Flexible override capability**: Component-specific settings allow for exceptions when needed, accommodating special requirements without compromising the overall update strategy

Technet24

You now see some of the many benefits to having a centralized update management system within the Veeam environment. Let's take a look at some best practices for update management.

Best practices for update management

To maximize the benefits of VSA 13's centralized update capabilities, consider the following recommendations:

- **Define appropriate maintenance windows**: Schedule update installations during periods of low backup activity. Consider time zones if your infrastructure spans multiple regions and use component-specific overrides where needed.
- **Monitor update status regularly**: Use the Veeam Updater interface accessible through the Web UI or Host Management console to verify that updates are being applied as expected across all components.
- **Test optional updates in non-production**: If your environment includes non-production VSA deployments, consider enabling automatic optional updates there first to validate changes before production deployment.
- **Plan for major releases**: Since major releases require manual action, establish a process for evaluating, testing, and deploying major version upgrades according to your organization's change management procedures.
- **Document mirror repository procedures**: For air-gapped environments, establish and document the process for synchronizing your local mirror repository to ensure updates are available when needed.
- **Leverage Veeam ONE integration**: If using Veeam ONE, take advantage of the update monitoring capabilities that provide alerts for update status and can even trigger automated failover in High Availability configurations when update-related health issues are detected.

The centralized update management and auto-update capabilities in VSA 13 represent a mature approach to maintaining the security and currency of backup infrastructure. By combining mandatory security updates with flexible optional update policies, Veeam has struck a balance between ensuring critical protection and respecting organizational change management requirements. This appliance-centric update model, managed through the Veeam Updater service, reduces administrative burden while improving security posture—a combination that makes VSA 13 particularly attractive for organizations seeking to modernize their data protection infrastructure.

Summary

This chapter has reviewed many of the new Veeam Software Appliance (VSA) features that are now in Veeam Backup and Replication v13. We looked at what the VSA is and what is new with this release. We covered how the VSA is security-hardened out of the box and supports a new High Availability (HA) mode. We reviewed the new Web UI for managing the VSA and the new Enhanced RBAC roles you can assign to users. Finally, we reviewed the central update management and auto-update features built into the VSA to keep it current.

Chapters 1 and *2* covered Veeam Backup and Replication basics for installation and upgrade, as well as the new Veeam Software Appliance (VSA). The next chapter, *Just Enough Operating System (JeOS)*, looks at the new way to deploy both Proxy servers and Veeam Hardened repositories in the Veeam Backup and Replication v13 infrastructure.

Further reading

- Veeam Software Appliance (VSA) Requirements - `https://helpcenter.veeam.com/docs/vbr/userguide/system_requirements.html?ver=13#backup-server`
- Veeam Windows to Appliance conversion request: `https://go.veeam.com/vsa-conversion`
- Creating custom roles - `https://helpcenter.veeam.com/docs/vbr/userguide/configure_roles.html?ver=13`
- Configuring Updates - `https://helpcenter.veeam.com/docs/vbr/userguide/update_appliance_configure_updates.html?ver=13`

Technet24

3

Just Enough Operating System (JeOS)

Just Enough Operating System (JeOS) is part of Veeam Backup and Replication v13, the latest release of Veeam components, and allows you to deploy new appliances, including proxies and **hardened repositories**. This chapter examines what JeOS is and how it is deployed from an ISO image. We will review the Veeam components you can create using JeOS, including **proxy servers** and hardened repository servers. We then examine how, just like the VSA in *Chapter 2, Veeam Software Appliance (VSA)*, JeOS is security-hardened out of the box. Lastly, we will examine the hardened repository and new immutability features such as block generation.

In this chapter, we're going to cover the following main topics:

- Veeam infrastructure appliance (JeOS)
- Veeam Proxy and hardened repository – JeOS Deployment
- Out of the box security hardening with JeOS ISO
- Hardened repository immutability features

Technical requirements

To ensure you can follow along with this chapter, you will require the following:

You must have downloaded the latest **ISO file** from `https://my.veeam.com/my-products`. You will have to register on the site and can obtain a trial license. As of this writing, version `13.0.1.1071` is the latest release.

Technet24

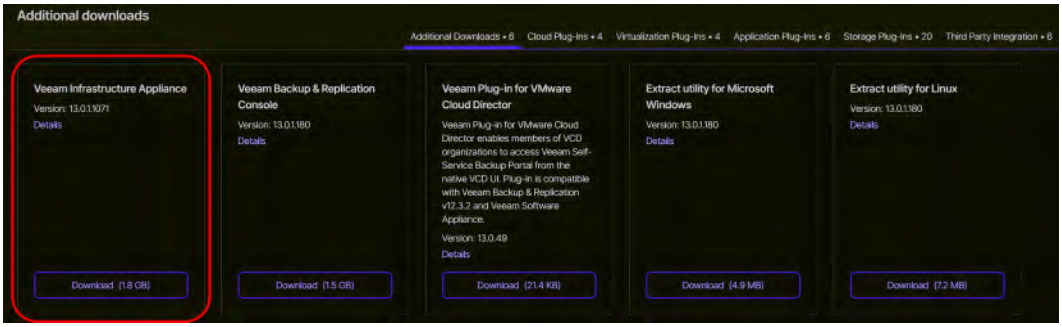


Figure 3.1: Additional downloads where the JeOS (VIA) ISO is located

Veeam infrastructure appliance (JeOS)

Veeam Backup and Replication v13 introduces a new deployment approach: the **Veeam Infrastructure Appliance (VIA)**, built on a purpose-designed operating system called **Just Enough Operating System** or **JeOS**. This minimalistic Linux distribution represents a fundamental shift in how organizations can deploy and manage Veeam backup infrastructure components, eliminating the traditional complexity of deploying and maintaining separate Linux or Windows servers for backup roles.

What is Veeam JeOS?

Veeam JeOS is a purpose-built, minimalistic Linux distribution developed and maintained entirely by Veeam to streamline the deployment and management of Veeam infrastructure components. Based on Rocky Linux, JeOS is delivered as a bootable ISO image and contains only the essential services and components required to operate specific Veeam roles. Unlike traditional deployment methods that require administrators to provision, harden, and maintain a full Linux or Windows operating system, JeOS provides a pre-hardened, security-optimized foundation specifically designed for Veeam workloads.

The JeOS philosophy follows a "*just enough*" principle: it includes only the essential services and components necessary for the intended backup infrastructure role. This approach dramatically reduces the attack surface, simplifies management, and ensures consistent deployment across your entire backup infrastructure.

JeOS vs Veeam software appliance

It is important to understand the distinction between JeOS and the Veeam Software Appliance (VSA). VSA is the complete backup server deployment that includes the JeOS operating system, the Veeam Backup and Replication or Veeam Backup Enterprise Manager server components, a PostgreSQL database, and a default hardened repository. The VSA is designed for deploying a complete backup server.

VIA, also known as the **JeOS ISO**, is a separate deployment option specifically designed for provisioning additional backup infrastructure components such as hardened repositories, backup proxies, WAN accelerators, and mount servers. Both the VSA and VIA share the same JeOS foundation but serve different purposes within the backup architecture.

Deployable infrastructure roles

The Veeam Infrastructure Appliance ISO supports deploying the following backup infrastructure roles:

Role	Description
Hardened Repository	Linux-based immutable repository with built-in ransomware protection. Eliminates SSH requirement and uses certificate-based authentication.
Infrastructure Appliance	Multi-purpose role for VMware backup proxies, general-purpose proxies, and WAN accelerators.
Infrastructure Appliance (with iSCSI & NVMe)	Extended appliance with additional components for Direct SAN proxy mode and NVMe-oF support.

Table 3.1: Infrastructure roles for deployment

Key benefits of JeOS deployment

JeOS deployment makes it easier to set up and get Veeam components installed quickly. Let's now explore the many benefits of the JeOS deployment.

- **Enhanced security posture**: JeOS deployment ensures enhanced security posture by considering the following:
    - **DISA STIG compliance:** The JeOS operating system is pre-hardened following Defense Information Systems Agency Security Technical Implementation Guides (DISA STIG), providing enterprise-grade security out of the box.
    - **Minimized attack surface:** By including only essential components, JeOS dramatically reduces potential vulnerability entry points compared to full operating system deployments.
    - **SSH disabled by default:** Secure Shell (SSH) access is disabled by default for security and hardening purposes and should only be enabled temporarily for troubleshooting purposes, eliminating a common attack vector.

Technet24

    - **Certificate-based authentication:** JeOS components use certificate-based authentication over secure, industry-standard communication protocols, eliminating the need for stored credentials.
    - **SELinux enforcing mode:** Security-enhanced Linux runs in enforcing mode, providing mandatory access control and preventing unauthorized system modifications.
- **Simplified operations**: JeOs deployment simplifies operations bt implementing the following:
    - **No Linux expertise required:** Administrators do not need deep Linux knowledge to deploy and manage JeOS-based infrastructure. The web-based management interface handles all configuration tasks.
    - **Centralized updates:** Veeam manages and distributes all OS and component updates through the Veeam Updater service, coordinated from the central backup server.
    - **Automated patching:** Updates can be scheduled during maintenance windows and applied automatically across all JeOS infrastructure components.
    - **Web-based host management:** The Veeam Host Management console (Cockpit-based) provides intuitive system administration without command-line interaction.
- **Rapid deployment**: JeOS can be deployed quickly because of the following:
    - **Quick installation:** Complete deployment typically takes under 10 minutes from boot to operational status.
    - **Guided configuration wizard:** Step-by-step wizard handles all initial configuration including networking, authentication, and security settings.
    - **Simplified pairing:** New JeOS appliances can be paired with the backup server using thumbprint verification and temporary PIN codes, eliminating complex credential management.
    - **Consistent deployments:** Every JeOS deployment follows the same hardened baseline, ensuring a consistent security posture across all infrastructure components.

Enhanced security, simplified operations, and rapid deployment are just some of the benefits for JeOS. We will now take a look at the system requirements for these deployments.

Deploying the Veeam infrastructure appliance

Before deploying the Veeam Infrastructure Appliance, ensure your target system meets the following minimum requirements:

System requirements

Table 3.2 lists the system requirements for JeOS deployments:

Component	Requirement
CPU	Minimum 2 cores (additional cores based on infrastructure role requirements)
Memory	Minimum 8 GB RAM for base OS (additional memory based on role requirements)
Storage	Minimum 2 disks: System disk (120 GB+) and Data disk (120 GB+)
Boot Mode	UEFI required (BIOS boot is blocked for security reasons)
Platform	Physical server, VMware ESXi 7.0U2+, Hyper-V, Proxmox VE, or other hypervisors
Hardware Compatibility	Red Hat Enterprise Linux 9 compatible hardware (RHEL HCL)

Table 3.2: System Requirements for JeOS deployments

The following sections detail the complete deployment process for deploying a JeOS-based infrastructure component. This example demonstrates deploying a hardened repository, though the process is similar for other infrastructure roles.

Step 1: Download the VIA ISO

Download the Veeam Infrastructure Appliance ISO from the Veeam Customer Portal. Navigate to **My Products** and locate the ISO under **Additional Downloads** | **Extensions and Other**. The ISO file is named *VeeamInfrastructureAppliance_13.0.1.1071_20251217.iso*. (Refer to *Figure 3.1*)

Step 2: Prepare the target system

Create a virtual machine or provision a physical server based on the system requirements. Ensure UEFI boot mode is configured and attach the ISO image as a boot device. For virtual deployments, configure at least two virtual disks: one for the operating system and one for data storage.

Technet24

Step 3: Boot and select role

Boot the system from the ISO. The boot menu presents you with the available installation options. Select the appropriate role for your deployment:

- **Veeam Hardened Repository**: For deploying an immutable backup repository.
- **Veeam Infrastructure Appliance:** For deploying proxies or WAN accelerators.
- **Veeam Infrastructure Appliance (with iSCSI & NVMe):** For Direct SAN or NVMe fabric deployments.

As shown below in *Figure 3.2*, we have the JeOS deployment screen on the first boot of the server.

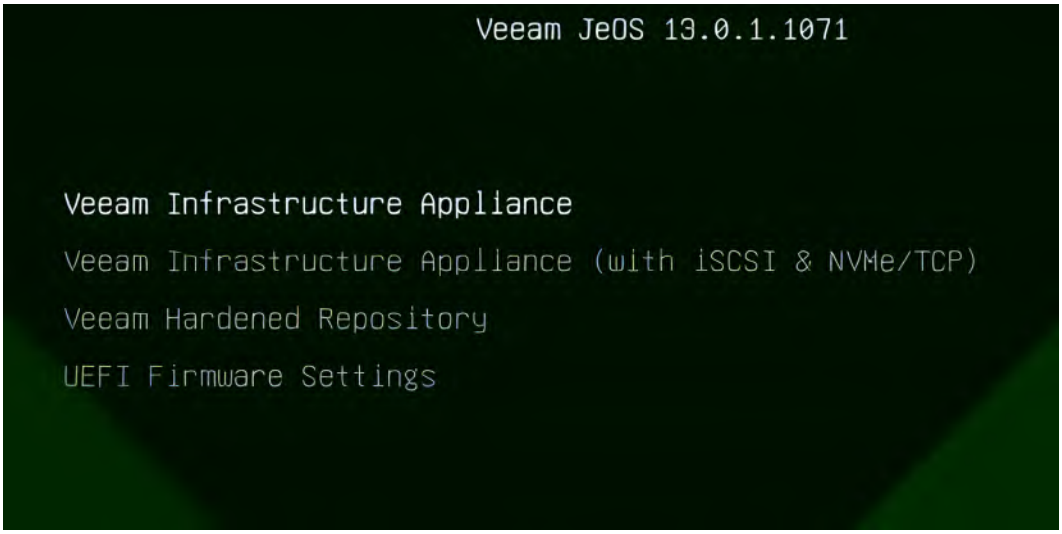


Figure 3.2: Selecting the role required for JeOS

As shown in *Figure 3.3*, you have the option to do a clean install which removes all existing data or reinstall if you need to save data such as backups on a repository server.

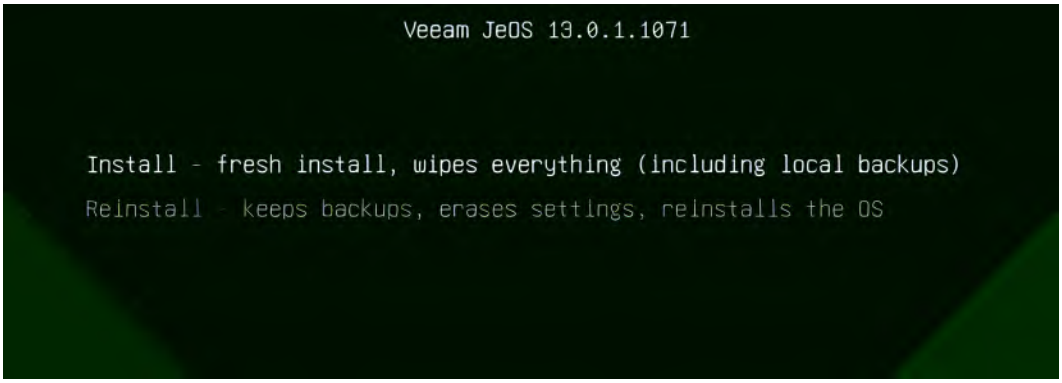


Figure 3.3: Install or reinstall options

Step 4: Automated installation

After selecting the installation type, the installer automatically checks disk requirements and proceeds with installation. A warning displays indicating that all existing data on the disks will be erased. The installation typically completes in 5-10 minutes. No internet connectivity is required during installation.

> **Note**
>
> During this step you cannot change anything while the installation runs.

In *Figure 3.4*, you can see the warning about wiping all the data with the **Install** option.

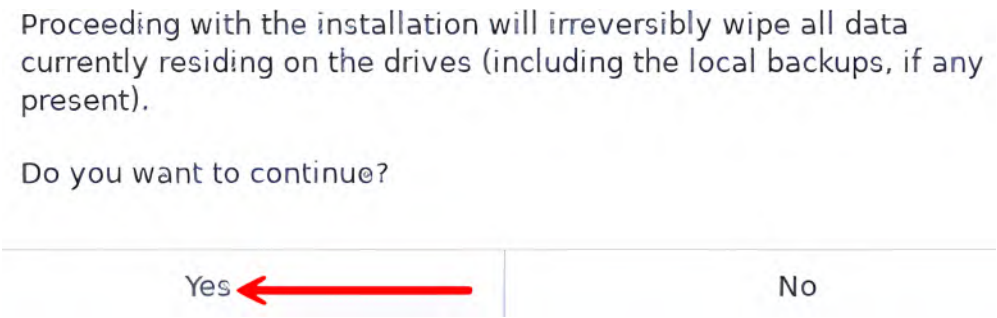


Figure 3.4: Notification about wiping all data

You can see the JeOS deployment progress screen in *Figure 3.5*.

Technet24

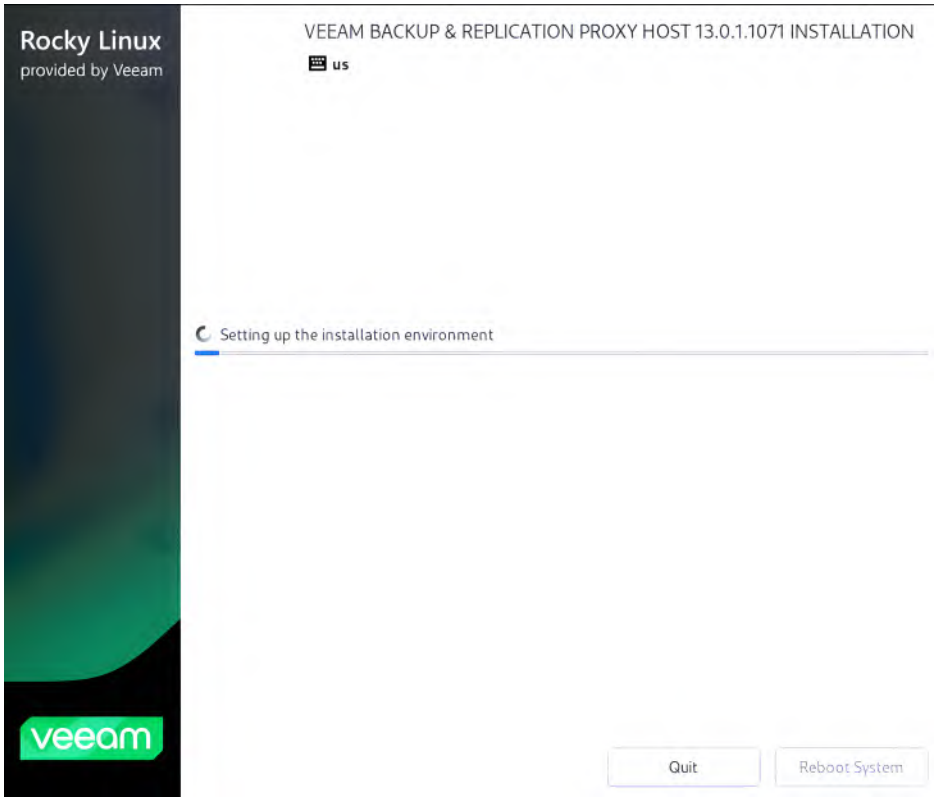


Figure 3.5: Installation of JeOS environment

Step 5: Initial configuration wizard

Following the reboot, the Initial Configuration Wizard launches. Complete the following configuration steps:

1. License agreement: Accept the **End User License Agreement** to proceed
2. Network configuration: Configure network settings, including IP address (static or DHCP), subnet mask, gateway, DNS servers, and hostname

3. Time configuration: Set the timezone and configure NTP server addresses for accurate time synchronization

> **Tip**
>
> Ensure to use servers that are reliable for NTP settings, as proper synchronization is required for both time and the MFA settings noted below in *Step 5*.

4. Host administrator account: Create the veeamadmin password. The password must comply with DISA STIG requirements (minimum 15 characters, complexity requirements). The password must comply with the following requirements:
    - 15 characters minimum.
    - 1 uppercase character.
    - 1 lowercase character.
    - 1 numeric character.
    - 1 special character.
    - No more than 4 characters of the same class in a row. For example, more than 4 lowercase or 4 numerical characters in sequence.
5. Multi-factor authentication: Configure MFA using a TOTP authenticator application such as Google Authenticator or Microsoft Authenticator. This is mandatory for Host Administrator accounts.
6. Security officer account (optional): Configure the veeamso Security Officer account for approving sensitive operations. While optional, this account is highly recommended for production deployments as it adds a second level of security for tasks that could compromise the system.

Step 6: Complete configuration

Review the configuration summary and click **Finish** to complete the setup. The system will finalize configuration and display the connection URLs for the Veeam Host Management console (port 10443) and any applicable web interfaces.

Technet24

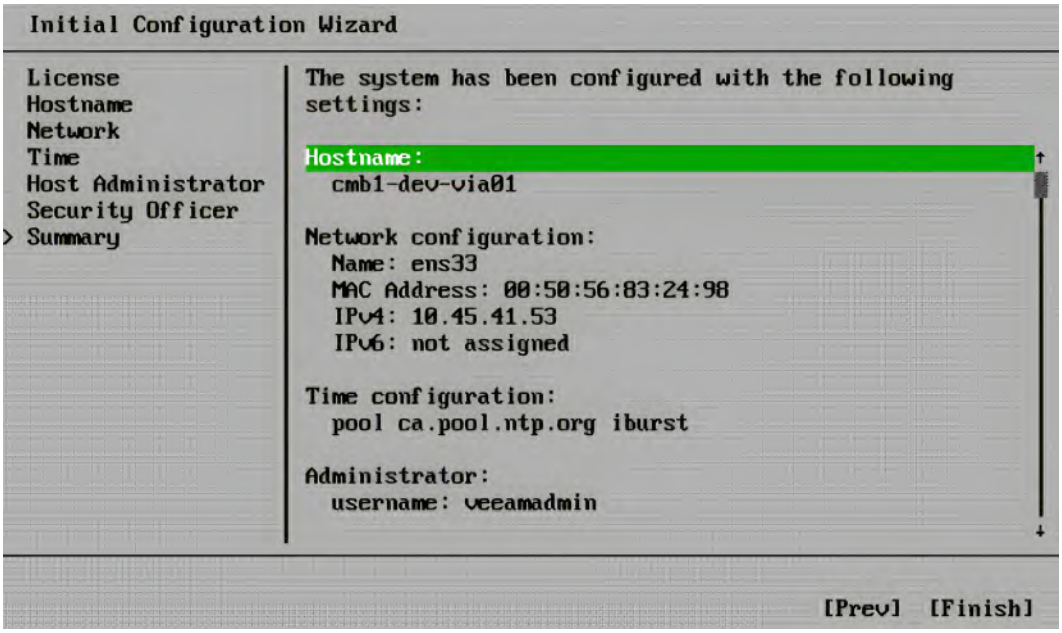


Figure 3.6: Summary screen before Finish

As shown in *Figure 3.6*, you have the JeOS Initial Configuration Wizard which is where you will enter information about your server including network settings and set up the administrator and security officer accounts.

Managing JeOS appliances

JeOS appliances are managed through the Veeam Host Management web interface, accessible via HTTPS on port 10443. This Cockpit-based interface provides comprehensive system administration capabilities without requiring command-line access.

Management Area	Capabilities
Network	IP configuration, hostname, DNS settings, proxy configuration
Time	Time zone configuration, NTP server management
Users and Roles	User account management, role assignment (Host Admin, User, Service Account)
Updates	View available updates, schedule update windows, apply patches
SSH	Temporarily enable/disable SSH for troubleshooting

Management Area	Capabilities
Applications	View paired backup servers, generate pairing codes, manage relationships
Security	Infrastructure lockdown, Security Officer approvals, password management

Table 3.3: Management areas for JeOS and capabilities

The Veeam Infrastructure Appliance represents a significant change in backup infrastructure deployment. By leveraging the JeOS foundation, organizations can deploy secure, hardened backup infrastructure components in minutes rather than hours or days. The elimination of traditional operating system management, combined with centralized updates and certificate-based authentication, dramatically reduces operational complexity and improves the security posture. For organizations seeking to modernize their backup infrastructure or expand existing Veeam deployments, the JeOS-based approach offers compelling benefits in security, simplicity, and operational efficiency.

Now, let us look at the components you can use with JeOS: *Proxies and hardened repositories*.

Veeam proxy and hardened repository – JeOS deployment

VIA provides a streamlined method for deploying essential backup infrastructure components including backup proxies and hardened repositories. Built on the same JeOS foundation as VSA, VIA delivers pre-hardened, security-optimized infrastructure components that can be deployed in minutes and managed through a web-based interface without requiring Linux expertise.

In this section, you will get detailed guidance on deploying both VMware backup proxies and Linux hardened repositories using the VIA JeOS ISO, covering the complete process from initial installation through integration with your Veeam Backup and Replication server.

Understanding VIA installation options

When booting from the VIA JeOS ISO, you are presented with three installation options at the GRUB boot menu. Each option includes a specific set of components tailored for different infrastructure roles:

Technet24

Installation Option	Use Case	Components Included
Infrastructure Appliance	VMware proxies, general-purpose proxies, WAN accelerators, mount servers	Minimal JeOS with Veeam transport services and data movers
Infrastructure Appliance (with iSCSI & NVMe)	Direct SAN proxy mode, NVMe fabric support, storage snapshots	Base components plus Open-iSCSI initiator and NVMe/TCP tools
Hardened Repository	Immutable backup storage with ransomware protection	Full security hardening, XFS with immutability support, Security Officer functionality

Table 3.4: Installation options for the VIA appliance

Note

Important:

The Infrastructure Appliance options can serve as standard Linux repositories but cannot function as hardened repositories with immutability. To deploy a hardened repository, you must select the **Hardened Repository** option during installation. This selection cannot be changed post-installation.

Part 1: deploying a VMware backup proxy

VMware backup proxies are essential components that retrieve data from production storage, compress and deduplicate it, and send it to the backup repository. Deploying proxies using VIA provides certificate-based authentication, automated updates, and eliminates the need to manage separate Linux or Windows servers for proxy roles.

System requirements for proxy deployment

The Proxy server has specific requirements, and the following table outlines what those requirements are:

Component	Requirement
CPU	Minimum 2 cores for base OS, plus 1 core (vCPU) for every 2 additional concurrent tasks. Using faster processors improves data processing performance.
Memory	Minimum 2 GB RAM, plus 1 GB for each concurrent task. The actual size of memory required may be larger and depends on the amount of data to back up, machine configuration, and job settings. Using faster memory improves data processing performance.
System Disk	Minimum 120 GB for operating system and Veeam components
Data Disk	Minimum 120 GB for logs, catalogs, and temporary data
Network	High bandwidth connection to vCenter/ESXi hosts and backup repository
Boot Mode	UEFI required (BIOS boot blocked for security)

Table 3.5: Proxy server deployment requirements

We have looked at the JeOS deployment options along with some of the system requirements, benefits, and installation steps. Now, we will go through the installation of the proxy role for the VIA.

Installing the infrastructure appliance for proxy role

The following steps will guide you through the installation process for the VIA appliance and the requirements for the Proxy Server role.

Note

Please refer to *Figures 3.2* to *3.6*, as the steps have the same options for the infrastructure appliance for the proxy role.

1. Download the VIA ISO: Access the **Veeam Customer Portal** and navigate to **My Products**. Locate the **Veeam Infrastructure Appliance ISO** under **Additional Downloads | Extensions and Other**. Download the **VeeamLiveOS** ISO file.

Technet24

2. Prepare the virtual machine: Create a new VM with UEFI boot mode enabled. Configure two virtual disks: a 120 GB system disk and a 120 GB data disk minimum. Allocate CPU and memory based on expected concurrent task count.
3. Boot from ISO: Attach the VIA ISO and boot the VM. From the GRUB menu, select **Veeam Infrastructure Appliance** (or **Veeam Infrastructure Appliance (with iSCSI & NVMe)** if Direct SAN access is required).
4. Automated installation: The installer validates disk requirements and proceeds automatically. Installation typically completes in 5–7 minutes. No internet connectivity is required.
5. Initial Configuration Wizard: After reboot, complete the configuration wizard:
    - Accept the **End User License Agreement**
    - Configure hostname (FQDN recommended).
    - Configure network settings (static IP recommended for infrastructure components).
    - Set time zone and NTP servers.
    - Create veeamadmin password (DISA STIG compliant: minimum 15 characters as noted in the previous section – step 4)
    - Configure MFA using a TOTP authenticator application.
    - Security officer account (optional): Configure the veeamso Security Officer account for approving sensitive operations. While optional, this account is highly recommended for production deployments as it adds a second level of security for tasks that could compromise the system.
6. Review and complete: Review the configuration summary and click **Finish**. Note the Veeam Host Management URL displayed (`https://[hostname]:10443`) so that you can connect to the management interface after the initial setup is completed.

Now that you've successfully installed the proxy server role, let's take a look at how you can add the proxy server to the Veeam environment.

Adding the VIA server to Veeam backup and replication

VIA-based infrastructure components use certificate-based authentication, eliminating the need for SSH or stored credentials. The initial connection uses a pairing code for secure trust establishment.

1. Open the Veeam Backup & Replication Console or Web UI and navigate to **Backup Infrastructure | Managed Servers**
2. Right-click on **Manage Servers** and select **Add Server | Linux**

3. Enter the DNS name or IP address of the VIA appliance
4. Next, select **Connect** using certificate-based authentication to establish an SSH connection
5. Verify the certificate fingerprint matches the appliance and confirm trust
6. For newly deployed appliances, a connection is automatically established with Veeam Backup & Replication. For previously paired appliances, generate a new pairing code from the **Host Management** console under **Applications**.

That takes care of adding the server to the Veeam environment. Let's now take a closer look at assigning the proxy role to the newly added server.

Assigning the VMware backup proxy role

Once the VIA server is added as a managed server, you can assign the VMware backup proxy role:

1. Navigate to **Backup Infrastructure** | **Backup Proxies**
2. Click **Add Proxy** and select **VMware backup proxy**
3. Next, to choose the server, select the **VIA server** from the **Choose server** dropdown
4. Configure **Transport mode** (Automatic selection is recommended for most deployments as this is the simplest method as well as the "Default" method used by Veeam)
5. Set **Max concurrent tasks** based on available CPU cores (best practice: 2 tasks per core)
6. Configure **Connected datastores** (Automatic detection or specify manually)
7. Configure traffic throttling rules if required
8. Review the summary and click **Apply** to complete proxy configuration

As you can see, adding the Proxy role is a very simple and straightforward process. The following table explains the supported transport modes for the Proxy role.

Supported transport modes

Table 3.6 illustrates the transport modes, requirements, and VIA installation options for Proxy server role:

Transport Mode	Requirements	VIA Installation Option
Network (NBD)	Network connectivity to ESXi hosts	Infrastructure Appliance

Technet24

Transport Mode	Requirements	VIA Installation Option
Virtual Appliance (HotAdd)	Proxy VM must have access to the datastore where the source VM resides	Infrastructure Appliance
Direct SAN (iSCSI)	iSCSI initiator configured, SAN fabric access	Infrastructure Appliance (with iSCSI & NVMe)
Direct SAN (FC)	Fibre Channel HBA, SAN fabric access	Infrastructure Appliance (with iSCSI & NVMe)
Direct NFS	NFS datastore access	Infrastructure Appliance

Table 3.6: Transport Modes for the Proxy Server Role

The preceding table looked at the transport modes for a proxy server; now let's discuss deploying a hardened repository in the next section.

Part 2: deploying a hardened repository

A hardened repository provides immutable backup storage that protects against ransomware, accidental deletion, and insider threats. When deployed using VIA, the hardened repository includes DISA STIG security hardening, an XFS filesystem with fast clone support, and Security Officer functionality for approving sensitive operations.

The following table outlines the system requirements for deploying a hardened repository.

System requirements for a hardened repository

In the following table, you can see the system requirements for a hardened repository:

Component	Requirement
Deployment type	Physical server strongly recommended (reduces attack surface)
CPU	The number of cores depends on the concurrent task settings. For more information, see *Limitation of Concurrent Tasks* (`https://helpcenter.veeam.com/docs/vbr/userguide/limiting_tasks.html#repo`)
Memory	Minimum 8 GB RAM (16 GB recommended for production)
System disk	Minimum 120 GB for operating system (SSD recommended)

Component	Requirement
Data disk(s)	Minimum 256 GB; locally attached block storage (SAS, internal drives)
Storage type	In addition to the standard *system requirements* (`https://helpcenter.veeam.com/docs/vbr/userguide/system_requirements.html`)for a backup infrastructure component, the machine where Veeam Infrastructure Appliance is installed must also meet the following requirements: [Hardened Repository only] Hardware RAID controllers must be used. Software RAID, Intel VMD VROC, and FakeRAID controllers are not supported
Boot mode	UEFI required with Secure Boot capability
Hardware compatibility	Red Hat Enterprise Linux 9 certified hardware (RHEL HCL)

Table 3.7: Hardened Repository system requirements

Note

Storage configuration:

The VIA installer automatically configures storage using LVM. The smallest disk is used for the operating system, while all remaining disks are combined into a single logical volume formatted with XFS for the backup repository. Plan your disk configuration accordingly before installation.

We will now look at the steps required to complete the Veeam hardened repository installation using the VIA appliance ISO.

Installing the hardened repository

To install the hardened repository, follow these steps:

1. Prepare the physical server: Configure RAID arrays as required. Ensure UEFI boot mode is enabled. Connect the server to the network but consider isolating it on a dedicated backup VLAN.
2. Boot from VIA ISO: Use a remote console (hardware vendor IPMI or out-of-band management) or bootable USB. In the GRUB menu, select **Veeam Hardened Repository**.

Technet24

3. Disk validation: The installer validates disk requirements. A warning indicates all existing data will be erased. Confirm to proceed with installation.
4. Automated installation: Installation completes in approximately 7–10 minutes. The system automatically partitions disks, installs the OS, and applies DISA STIG hardening.
5. Initial configuration wizard: After reboot, complete the following mandatory configuration steps:
    - License Agreement: Accept the EULA
    - Network configuration: Configure static IP address, subnet, gateway, and DNS servers
    - Time configuration: Set time zone and NTP servers (critical for immutability timestamps)
    - Host administrator (veeamadmin): Create a password meeting DISA STIG requirements
    - MFA configuration: Configure TOTP authentication (mandatory for hardened repository)
    - Security Officer (veeamso): Required for hardened repository; create password and configure MFA
6. Complete setup: Review configuration and click Finish. The system displays the Host Management URL for browser-based administration.

Now that you've installed the hardened repository, we'll move on to adding it to the Veeam environment.

Adding the hardened repository to veeam backup and replication

Follow these steps to add the hardened repository to the Veeam environment. This is similar to the process of adding a Proxy server, which we discussed in the previous sections.

Note

Security Officer Account:

The Security Officer (*veeamso*) is required for hardened repository deployments. This account approves sensitive operations such as removing immutability, extending immutability periods, or re-pairing with backup servers. Store the Security Officer

> credentials separately from regular administrator credentials to maintain separation of duties.

1. In the VBR Console or Web UI, navigate to **Backup Infrastructure** | **Managed Servers**
2. Right-click and select **Add Server** | **Linux**
3. Enter the hardened repository hostname or IP address
4. Select **Connect using certificate-based authentication**
5. Verify the certificate fingerprint and establish trust
6. For newly deployed appliances, a connection is automatically established with Veeam Backup & Replication. For previously paired appliances, generate a new pairing code from the **Host Management** console under **Applications**

We will now add the repository and configure immutability for the hardened repository.

Configuring the backup repository with immutability

Follow these steps to configure the backup repository with immutability:

1. Navigate to **Backup Infrastructure** | **Backup Repositories**
2. Click **Add Repository** and select **Direct Attached Storage** | **Linux (Hardened Repository)**
3. Provide a descriptive name and optional description for the repository
4. Select the VIA hardened repository server and click **Populate** to discover available paths
5. Select the repository path (typically `/var/lib/veeam` or the default mount point)
6. Click **Populate** to verify capacity and free space
7. Enable **Use fast cloning** on XFS volumes to leverage block cloning for synthetic full operations
8. Enable **Make recent backups immutable** for [X] days and specify the immutability period
9. Configure load control settings to limit concurrent tasks based on server capacity
10. Select **Mount Servers for Windows and Linux restore operations**
11. Review the configuration summary and click **Apply** to create the repository

With this, you complete the steps needed to deploy and add a hardened repository to the Veeam infrastructure. Now let's look at some of the considerations for immutability.

Technet24

Immutability considerations

When configuring immutability for a hardened repository, consider the following important factors:

- **Backup method restriction:** Only forward incremental backup methods with active full or synthetic full are supported. Reverse incremental and forever forward incremental methods cannot be used with immutable repositories.

> **Note**
>
> With version 13 of Veeam Backup and Replication, the Reverse Incremental backup method has been deprecated so it is no longer an option to select.

- **Retention alignment:** The immutability period should align with your backup retention policy. Files cannot be deleted or merged until immutability expires.
- **Backup copy jobs:** To use immutability with backup copy jobs, you must enable the GFS (Grandfather-Father-Son) retention policy.
- **Metadata files:** VBM metadata files cannot be immutable as they are updated with each backup session. Veeam implements additional protections for these files.
- **Capacity planning:** Plan storage capacity carefully as immutable files cannot be removed before expiration, even if the repository approaches capacity.

The Veeam Infrastructure Appliance provides a streamlined, security-focused approach to deploying essential backup infrastructure components. By leveraging the VIA JeOS ISO, organizations can deploy both backup proxies and hardened repositories with consistent security baselines, certificate-based authentication, and centralized update management. The key distinctions between deployment options ensure that administrators select the appropriate installation type based on their specific infrastructure requirements, whether deploying high-performance proxies for data movement or immutable repositories for ransomware protection.

Quick reference: VIA deployment summary

Aspect	Proxy Deployment	Hardened Repository
Installation Option	Infrastructure Appliance	Hardened Repository
Recommended Platform	Virtual or Physical	Physical (strongly recommended)

Aspect	Proxy Deployment	Hardened Repository
Security Officer	Optional	Required
Authentication	Certificate-based	Certificate-based
Storage Requirement	Local or network-attached	Direct-attached block storage only
Immutability Support	No (standard repository only)	Yes (XFS with fast clone)
MFA requirement	Host Admin only	Host Admin and Security Officer

Table 3.8: VIA appliance deployment summary

Now that you've learned how to deploy the components of the VIA, namely proxy servers and the hardened repository, let's discuss how the VIA is hardened out of the box for security.

Out-of-the-box security hardening with JeOS ISO

One of the most significant security advancements in Veeam Data Platform v13 is the introduction of the JeOS ISO. This purpose-built, minimalistic Linux distribution represents a fundamental shift in how backup infrastructure can be deployed and secured. Rather than requiring administrators to manually harden operating systems or rely on third-party security configurations, the JeOS ISO provides a pre-hardened foundation that meets enterprise security standards from installation.

The JeOS ISO is developed and maintained entirely by Veeam, eliminating the traditional overhead associated with securing Linux repositories while providing a consistent, validated security baseline across the entire backup infrastructure.

DISA STIG compliance by default

The JeOS ISO is hardened to DISA STIG standards out of the box. DISA STIGs are among the most rigorous security frameworks available, originally developed for U.S. Department of Defense systems and widely adopted across government agencies and security-conscious enterprises worldwide.

Prior to v13, achieving DISA STIG compliance for Veeam infrastructure required administrators to manually execute hardening scripts, configure audit rules, implement password policies, and validate hundreds of individual security settings. This process was time-consuming, error-

Technet24

prone, and required significant Linux expertise. JeOS ISO eliminates this burden entirely by shipping with all applicable DISA STIG requirements preconfigured.

Key DISA STIG categories implemented

JeOS ISO implements security controls across several critical categories:

- **Password policies:** Complex password requirements are enforced for all user accounts, including minimum length, character complexity, and password history restrictions. Passwords must meet DISA STIG-compliant complexity standards during initial configuration.
- **Audit and logging:** Comprehensive audit rules configured with `auditd` to track account creation, modification, and termination events. File integrity monitoring is enabled through AIDE to detect unauthorized changes to critical system files.
- **Service minimization:** Insecure legacy services such as `telnet` and `rsh-server` are removed. Only essential services required for backup operations are present, dramatically reducing the attack surface.
- **TCP/IP hardening:** Network stack security settings are configured to prevent common network-based attacks, including SYN flood protection and ICMP hardening.
- **File system permissions:** Restrictive permissions are applied to audit logs, configuration files, and sensitive system directories to prevent unauthorized access.
- **Session management:** Session locking capabilities are enabled with `vlock`, allowing administrators to secure console sessions when stepping away.

With these key security categories implemented by default, it is ensured that your environment is secure. Now, let's look at the Zero Trust architecture for the VIA appliance.

Zero trust architecture and privilege separation

JeOS ISO implements a true Zero Trust architecture that fundamentally changes how backup administrators interact with the underlying system. In traditional deployments, backup administrators typically have root or elevated privileges on the backup server, creating a significant security risk. If an attacker compromises administrator credentials, they gain full control of the backup infrastructure.

With JeOS-based Veeam Software Appliance, backup server administrators do not receive root access or base OS privileges. All management tasks are performed through the Web UI or the dedicated Host Management console on port 10443. This separation ensures that even if backup administrator credentials are compromised, the attacker cannot escalate to operating system-level access.

Designated security officer role

JeOS architecture introduces a Designated Security Officer who must approve sensitive management tasks. This four-eyes principle is where one person initiates a task/change and the designated security officer reviews and approves it. This ensures that critical infrastructure changes require explicit authorization from a separate security role, preventing single points of compromise. When combined with mandatory multi-factor authentication, this creates multiple layers of protection around administrative operations.

Built-in immutability and ransomware protection

Every deployment using JeOS ISO includes a built-in hardened repository with immutability capabilities. This immutability is not a separate configuration option but an integral part of the appliance architecture. Backups stored in the default repository can be protected with configurable retention periods that prevent any user, including administrators, from modifying, deleting, or encrypting them.

The Veeam Hardened Repository component has passed external audits for compliance with the Write Once Read Many (WORM) storage standard, meeting the highest regulatory standards for data immutability. This certification provides assurance that the immutability implementation is technically sound and legally defensible for compliance purposes.

Backup infrastructure lockdown mode

JeOS-based appliance supports **Backup Infrastructure Lockdown** mode, which prevents unauthorized additions of backup infrastructure components. Once enabled, new proxies, repositories, or other infrastructure elements cannot be added without explicit approval from the Security Officer. This prevents an attacker who gains backup administrator access from adding rogue components to exfiltrate or manipulate backup data.

Automated security updates with veeam updater

One of the most significant operational benefits of JeOS ISO is the integration of Veeam updater for automated patching. Traditional Linux-based backup infrastructure requires administrators to manage operating system updates separately from Veeam software updates, often leading to delayed patches and security vulnerabilities.

Veeam Updater provides centralized, orchestrated updates for the entire appliance stack: the Veeam Data Platform software, the underlying JeOS operating system, and all managed components, including proxies and repositories, deployed from the same JeOS ISO. This unified approach ensures a consistent security posture across the entire backup infrastructure without requiring manual intervention or specialized Linux expertise.

Technet24

Update scheduling and control

Updates can be scheduled during maintenance windows to minimize operational disruption. Administrators retain control over when updates are applied while benefiting from Veeam-validated patches tested for compatibility with the backup infrastructure. This eliminates the common concern that operating system updates will break backup functionality.

Minimized attack surface

"*Just Enough OS*" is not merely a marketing term but a deliberate design philosophy. A JeOS distribution includes only the essential services and components required for backup operations. There is no graphical user interface, no unnecessary network services, and no additional packages that could introduce vulnerabilities.

SSH disabled by default

By default, SSH is disabled on JeOS deployments to follow security best practices. While SSH can be temporarily enabled for troubleshooting purposes through the console interface, the default configuration blocks this common attack vector. All routine management tasks are performed through the Web UI and Host Management console, eliminating the need for direct shell access during normal operations.

Security benefits — summary

The following table summarizes the key security improvements that the JeOS ISO provides compared to traditional manual deployments:

Security Area	Traditional Deployment	JeOS ISO Deployment
DISA STIG compliance	Manual script execution and validation required	Pre-configured and validated out of the box
OS patching	Administrator responsibility; often delayed	Automated via Veeam Updater
Root Access	Available to backup administrators	Not available; Zero Trust architecture
Attack Surface	Full Linux distribution with unused services	Minimal services; only back up essentials
SSH Access	Typically enabled for management	Disabled by default; console-based management

Security Area	Traditional Deployment	JeOS ISO Deployment
Immutability	Requires separate configuration	Built-in with WORM certification
Security Updates	Separate OS and Veeam update cycles	Unified, orchestrated updates
Compliance Evidence	Manual documentation required	Standardized, auditable configuration

Table 3.9: Security benefits of VIA appliance deployment

Now that we have looked at the many security features and benefits for the VIA appliance, let's move on to the operational implications that could come from this type of deployment.

Operational implications

The security hardening built into JeOS ISO addresses one of the most significant challenges in enterprise backup deployments: ensuring consistent security configuration across infrastructure that is often deployed by administrators without deep Linux security expertise. By shifting security responsibility from individual administrators to the Veeam engineering team, organizations benefit from validated security configurations that remain up to date through automated updates.

For compliance-focused organizations, the pre-hardened JeOS ISO simplifies audit processes. Rather than documenting and validating hundreds of individual security settings, auditors can reference the DISA STIG compliance built into the platform. The external WORM certification for the hardened repository provides additional audit evidence for data immutability requirements.

The combination of DISA STIG hardening, Zero Trust architecture, built-in immutability, automated patching, and minimized attack surface makes JeOS ISO the most secure foundation available for Veeam backup infrastructure. If your organization is deploying new backup environments or refreshing existing infrastructure, you should strongly consider JeOS-based Veeam Software Appliance as your primary deployment model.

We will now turn to our final topic: the hardened repository with new immutability features, including block generation.

Hardened repository immutability features

VIA hardened repository represents Veeam's most secure on-premises backup storage solution. Built on a pre-hardened Rocky Linux operating system compliant with DISA STIG security

Technet24

standards, the VIA Hardened Repository provides robust protection against ransomware, malicious insiders, and accidental data deletion through advanced immutability features.

Immutability is the state of data that prevents it from being modified or deleted for a specified period. When you configure a hardened repository, you specify the immutability period during which backup files cannot be moved, modified, or deleted, though they can still be copied for restore operations.

The VIA Hardened Repository leverages XFS's native immutable file attribute. For each backup file, Veeam Backup and Replication creates a corresponding `.veeam.N.lock` file containing immutability period information. Additionally, the `xattr` (extended attribute) with the immutable time period is set directly on each backup file, providing dual-layer protection.

How immutability works

The Veeam Immutability Service (*veeamimmureposvc*) runs with root privileges on the hardened repository and operates as a local service with no network access. This service performs the following critical functions:

- Checks immutability attributes every 20 minutes
- Calculates immutability time for files and sets or removes the immutable attribute
- Performs time drift detection to prevent clock manipulation attacks
- Maintains the `/etc/veeam/immureposvc/timeLog` file, updated every 10 minutes with UTC timestamps

This architecture prevents backup files from being deleted or modified even if an attacker compromises the Veeam Data Mover Service, Veeam Transport Service, or the NTP server. The immutability service operates independently and cannot be influenced by external network connections.

Block generation: optimizing immutable operations

Block generation is an automatic optimization mechanism that extends the immutability expiration date by up to 10 days (or 30 days for Amazon S3 and IBM Cloud object storage). This feature reduces I/O operations and associated costs without requiring any manual configuration.

How block generation works

Consider a scenario where you configure an immutability period of 30 days for your hardened repository:

1. When the first full backup arrives, its immutability period is set to 30 + 10 = 40 days. This full backup starts the first "generation" of backups.

2. All incremental backups created within the 10-day block generation period share the same immutability expiration date as the full backup. A data block transferred on day 9 will have the same expiration date as one transferred on day 1.
3. On the 11th day, a new block generation period begins. The immutability period must then be extended for dependent data blocks to ensure all blocks remain protected for the configured minimum period.
4. Reused data blocks are automatically assigned to new generations, continuously extending their immutability expiration period as long as they remain part of an active backup chain. Please refer to this link for more on the block generation process and diagrams - `https://helpcenter.veeam.com/docs/vbr/userguide/object_storage_block_generation.html?ver=13`

Block generation is a straightforward process if you understand the formula used for it. Now we will look at the benefits of the new block generation.

Benefits of block generation

The block generation mechanism provides several significant advantages:

- Reduced I/O operations: By batching immutability extensions within generation periods, Veeam dramatically reduces the number of attribute update operations required.
- Lower storage costs: For object storage repositories, fewer API calls translate directly to reduced transaction costs from cloud providers.
- Backup chain integrity: Within a forward incremental backup chain, incremental backup files depend on the full backup. Block generation ensures all dependent blocks maintain consistent immutability protection.
- Automatic operation: Block generation requires no configuration—it activates automatically when immutability is enabled, simplifying administration.

Block generation periods by storage type

The following table summarizes the default block generation periods for different storage types:

Storage Type	Block Generation Period
Amazon S3 Object Storage	30 days
IBM Cloud Object Storage	30 days

Technet24

Storage Type	Block Generation Period
Google Cloud Storage	30 days
Azure Blob Storage	10 days
S3-Compatible Storage	10 days
Hardened Repository (on-premises)	Up to 10 days

Table 3.10: Block generation periods based on storage type

Planning storage capacity with block generation

When planning storage capacity for immutable repositories, you must account for the additional retention period introduced by block generation. The formula for calculating required storage capacity is:

Total Retention = Job Retention + Immutability Period + Block Generation Period

For example, if you configure:

- Backup job retention: 30 days
- Immutability period: 14 days
- Block generation: 10 days (automatic)

You should plan storage capacity to retain backups for 30 + 14 + 10 = 54 days to ensure all immutable data can be stored without capacity constraints.

XFS fast clone and block clone integration

The VIA Hardened Repository uses the XFS file system specifically because it supports both immutable file attributes and the reflink (reference link) technology that enables Veeam's Fast Clone feature. Fast Clone, also known as block cloning, provides copy-on-write functionality that dramatically improves performance and storage efficiency.

How fast clone works

When Fast Clone is enabled on an XFS volume, Veeam uses the reflink capability to create lightweight copies of data blocks. Instead of physically copying data, the file system creates references to existing blocks. This technology is used in several key operations:

- **Synthetic full backups**: Creating synthetic full backups from incremental chains completes in seconds rather than hours, as only metadata needs to be updated

- **GFS backup creation**: Grandfather-Father-Son retention backups leverage block cloning, consuming virtually no additional storage space while maintaining separate immutability periods
- **Backup file merging**: When retention policies require merging backup files, Fast Clone performs the operation at the metadata level
- **Full backup file compacting**: Compacting operations to reclaim space from deleted data complete rapidly with minimal I/O overhead

These are the key operations performed by the fast clone process, which makes backup generation more performant. Let's move on to discuss the requirements for using XFS to enable the fast cloning technology.

XFS configuration requirements

For Fast Clone to function correctly on the VIA Hardened Repository, the XFS file system must be configured with specific parameters:

- Reflink must be enabled (reflink=1)
- Cyclic Redundancy Check (CRC) must be enabled (crc=1)
- Block size must be between 1 KB and 4 KB (4096 bytes recommended)
- The Linux kernel must support reflinks (all kernels in VIA meet this requirement)

The VIA pre-configures XFS with these optimal settings during initial deployment, ensuring Fast Clone capabilities are immediately available without manual configuration. The Veeam JeOS handles all file system configuration automatically.

GFS backups and immutability

Grandfather-Father-Son (GFS) backups receive special treatment in the immutability framework. When GFS retention is enabled, Veeam automatically makes GFS backup files immutable for the entire duration of their retention policy—not just the repository's configured immutability period.

For example, if you configure a yearly GFS backup with a 7-year retention policy and your repository immutability is set to 30 days, the GFS backup remains immutable for the full 7 years. This eliminates the need to manually align immutability periods with GFS retention policies.

Because Fast Clone creates GFS backups using block references rather than physical copies, organizations can maintain multiple immutable GFS copies (weekly, monthly, yearly) without significantly increasing storage consumption. The shared data blocks maintain immutability across all referencing backup files.

Technet24

Best practices for immutability configurations

When configuring immutability on your VIA Hardened Repository, consider the following recommendations:

- **Set appropriate immutability periods**: The minimum immutability period is 7 days, with a maximum of 9,999 days. Choose a period that provides sufficient time to detect and respond to ransomware attacks, typically 14 to 30 days for most organizations.
- **Use forward incremental backup methods**: Immutable repositories require forward incremental backup methods with active full or synthetic full backups enabled. Reverse incremental and forever forward incremental methods are not supported because immutable files cannot be merged or deleted until their immutability period expires.
- **Enable GFS for backup copy jobs**: Backup copy jobs targeting hardened repositories require GFS retention policy to be enabled. This ensures proper immutability handling for copied backups and prevents warnings during job configuration.
- **Account for block generation in capacity planning**: Always include the block generation period (up to 10 additional days) when calculating storage requirements. This prevents capacity issues when immutable data cannot be removed as expected.
- **Monitor time drift detection**: If the moveTime or accelerationTime parameter in the timeLog file exceeds 86,400 seconds (24 hours), retention operations will be blocked with a warning. This protection prevents attackers from manipulating system time to bypass immutability. Recovery requires root-level access to the hardened repository.

The VIA Hardened Repository's immutability features provide comprehensive protection for backup data through multiple integrated mechanisms. The combination of XFS file system immutability attributes, the isolated Veeam Immutability Service, automatic block generation optimization, and Fast Clone technology creates a robust defense against ransomware and malicious data manipulation.

By understanding how these features work together, particularly the automatic block generation that optimizes I/O operations while maintaining minimum immutability periods, administrators can properly plan storage capacity and configure retention policies that meet both security and operational requirements.

Summary

This chapter has provided you with an understanding of what VIA JeOS is for Veeam Backup and Replication v13, as well as the components that comprise the installation. We discussed what JeOS is and how you deploy it within your infrastructure. Then, we looked at the components which you can deploy from the JeOS, those being a Proxy Server or Veeam

Hardened Repository. Afterward, we examined how JeOS deployments are security-hardened from the moment they are deployed. Lastly, we examined the new immutability features, including the new block generation option. In the next chapter, we will cover the new Veeam Console for Veeam Backup and Replication v13.

Further reading

- Veeam Infrastructure Appliance (VIA) Requirements - `https://helpcenter.veeam.com/docs/vbr/userguide/system_requirements.html?ver=13#veeam-infrastructure-appliance`
- VMware Backup Proxy Requirements - `https://helpcenter.veeam.com/docs/vbr/userguide/system_requirements.html?ver=13#vmware-backup-proxy`
- Hardened Repository Requirements - `https://helpcenter.veeam.com/docs/vbr/userguide/system_requirements.html?ver=13#backup-repository`

Technet24

4

Veeam Backup & Replication Console

Veeam Backup & Replication console is part of Veeam Backup & Replication v13, the latest release of Veeam components, and is used in the management of your backup environment from jobs and repositories to tape servers and storage integration. This chapter examines what the Veeam Backup & Replication console is, including a brief history. We will review the deployment and navigation of the Veeam Backup & Replication console, examine updating the console and look at the differences between the web UI and Windows UI. Lastly, we will examine some advanced tips and tricks for using the Veeam console.

In this chapter, we're going to cover the following main topics:

- History of Veeam Backup & Replication user interface
- Deployment of the Veeam console
- Navigating the Veeam console
- Updating the Veeam Backup & Replication user interface
- Web UI vs. Windows UI in V13
- Advanced tips and tricks for using Veeam Console

History of Veeam Backup & Replication user interface

The **Veeam Backup & Replication user interface** (**UI**) for Windows environments has a surprisingly rich history—one that helps explain its longevity, consistency, and usability today. Veeam's origins trace back to the early 2000s, when VMware environments faced a fundamental challenge: accessing and moving files in and out of the VMFS file system. Veeam

Technet24

addressed this gap with FastSCP, a lightweight yet extremely fast file transfer utility for VMware. At the time, this capability alone set Veeam apart and drove much of its early adoption.

As the product evolved, FastSCP was eventually deprecated and fully integrated into what became the Veeam Backup & Replication platform. Over time, the UI expanded far beyond its original purpose, incorporating backup, replication, file copy, file-level restore, unstructured data protection, application-aware backups, and support for Windows, Linux, and macOS agents alongside an ever-growing list of hypervisors and platforms. What began as a simple utility matured into a comprehensive data protection console.

One of the most noteworthy aspects of the Veeam UI is its remarkable consistency. Administrators who used early versions from 2008, 2009, or 2010 will recognize familiar workflows even in modern releases. This continuity is intentional and reflects a strong emphasis on usability, training efficiency, and minimizing disruption across upgrades. From a design perspective, Veeam has long prioritized intuitive, wizard-driven workflows. A commonly cited rule of thumb is that nearly every task can be completed in seven steps or fewer—a principle that significantly reduces complexity and lowers the skills barrier. Paired with clear, embedded help content, and robust documentation, the UI enables administrators to quickly understand job organization, configuration options, and operational status without extensive ramp-up time.

This focus on ease of use is core to the company's founding philosophy. Early design goals stated that a new user should be able to download the product, install it, and configure their first backup in under ten minutes, all without consulting documentation. That principle continues to influence the UI design today and is a key reason for Veeam's broad adoption.

Over the years, the UI has undergone incremental improvements, with occasional major refreshes. One significant update occurred in the mid-2010s, aligning closely with the Microsoft ecosystem, which at the time played a central role in enterprise IT strategies. While Microsoft technologies remain important, the industry has since shifted toward cloud platforms, SaaS services, and more heterogeneous environments. As Veeam itself evolved into a broader platform company, the UI adapted accordingly by supporting new deployment models, architectures, and use cases.

With Veeam Backup & Replication v13, the Windows UI received what is likely its final major update. From this point forward, changes are expected to be more incremental, as strategic investment increasingly shifts toward the web-based user interface. The Windows UI remains fully supported and functionally complete, and organizations deploying it today can reasonably expect to operate it across an entire infrastructure lifecycle. At present, administrators should expect to work with both interfaces. While the web UI continues to expand in capability, some advanced features still require the Windows console. The long-term

development goal, however, is full feature parity in the web UI, enabling a more flexible and modern administrative experience.

Looking ahead, organizations will have choices of deploying Veeam Backup and Replication on Windows, adopting the Veeam software appliance, or transitioning between models as capabilities mature. Migration paths are planned, and most core features are available across both form factors. Throughout v13 it is reasonable to expect that more and more capabilities will come to the web user interface; however, the Windows UI will still be very important for day-to-day administration tasks. For example, the next major release will bring new capabilities into the web user interface, but not all of them yet. A good example in other technology ecosystems is how VMware's vCenter Server application for a while required use of both the Windows application and the web interface.

In summary, the Veeam Backup & Replication user interface represents a careful balance between innovation and familiarity. Its evolution reflects changing infrastructure realities, while its consistency underscores a deep commitment to usability, ensuring that administrators can adopt new capabilities without re-learning the product from scratch.

Now that we have familiarized ourselves with a brief history of the Veeam console, let's look at how you go about deploying the Veeam console.

Deployment of the Veeam console

Starting with Veeam Backup & Replication version 13, installation of the **Windows-based backup console** begins by downloading the installer from the web interface associated with the Veeam Backup & Replication server. This approach applies to both deployment models: the **VSA** and the **installable Windows application**.

The key difference between the two lies in identity integration. The Windows-based application is typically integrated with **Active Directory**, whereas the VSA is **not integrated by default**, although it can be configured to do so. *Figure 4.1* illustrates how the Veeam Backup & Replication console can be downloaded directly from the web UI, accessed via the default HTTPS address or server name of the Veeam Backup & Replication server.

Technet24

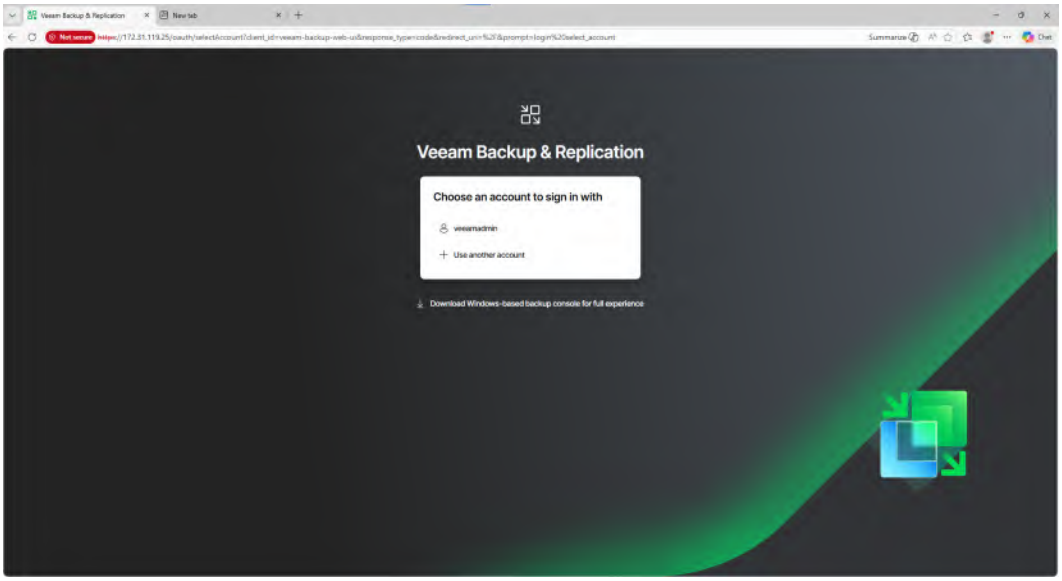


Figure 4.1: Download Windows-based console link on the web UI page

From the initial release of version 13, the standalone Windows-based backup console requires approximately 1.5 GB of disk space, so it is important to plan accordingly. In many environments, administrators may prefer to host the installer centrally rather than having large installation files scattered across individual download folders.

It is also worth noting that when the Veeam Backup & Replication server is updated, the backup console typically requires a corresponding update. These updates are automatically prompted when the console connects to an upgraded server, making the process straightforward and largely hands-off.

If Veeam Backup & Replication is installed on a Windows system, the backup console is included by default. However, administrators often connect from remote systems, such as a **jump box**, in which case installing the console on a separate system is recommended. At this point, ensure that day-to-day administration is not using the main production AD identity. It's a good idea to use a management fabric for Veeam operations to protect against any security risks.

Finally, in environments that use a mix of Windows-based Veeam Backup & Replication servers and the VSA, administrators must explicitly connect to the software appliance from a system that is ***not*** running the Windows-based Veeam Backup & Replication server. For mixed environments, using a dedicated jump box for day-to-day administration is strongly recommended.

Navigating the Veeam console

The first step in working with the Veeam Backup & Replication user interface is simply logging in. Once connected, you are presented with the **Home** view, which serves as the starting point for most interactions. Taking time to orient yourself within this view is important, as it provides a clear overview of the environment and the primary capabilities of the product.

At the top of the interface, a ribbon of icons provides quick access to common actions such as creating jobs, performing recoveries, and so on, as well as high-level environment information. On the left side of the console, the interface is logically organized into sections such as **Primary Jobs**, **Secondary Jobs**, **Restore**, and **Actions**. This structure makes it easy to understand what the Backup & Replication console can do and where to begin for common administrative tasks.

On the top-right corner of the interface, you will find the **Veeam Intelligence** widget displayed in a compact but prominent format. This AI-powered feature allows administrators to ask natural-language questions and quickly retrieve insights about the Veeam Backup & Replication environment, making it a helpful starting point for exploration and troubleshooting.

Also on the left side of the **Home** view are the primary navigation options, including **Jobs**, **Backups**, and **History**. Selecting any of these presents a corresponding inventory in the main pane on the right, offering a clear and centralized view of objects and activities. Along the lower left portion of the interface, additional functional areas are listed, which allows you to explore and enable capabilities as needed.

The user interface is also highly customizable. By selecting the ellipsis (three dots) menu, you can choose to show or hide specific sections. For example, if tape functionality is not used in an environment, it can be hidden to simplify the interface and reduce visual clutter.

Finally, attention should be given to the bottom-right corner of the console. This area displays critical contextual information at a glance, including the following:

- Currently logged-in user
- Connected Veeam server
- Product build number
- License information

This information is especially valuable when verifying connectivity or ensuring the environment is running a current build, which is a best practice for any production system. It is also a good indicator of consistency across environments with many Veeam Backup & Replication servers.

Technet24

Let's look at each of the sections within the Veeam console and how you put these capabilities to work.

Home section

The home section is the little house icon and is the starting point for all features in the backup console. The main areas in this section are the **Main Menu**, **Jobs View**, **Backups View**, and **Last 24 Hours View**. Let's explore these areas in detail.

Main menu

Many of the options in the main menu console are self-explanatory, but I will describe some of the highlights. *Figure 4.2* shows the view of the main menu for a V13 server:

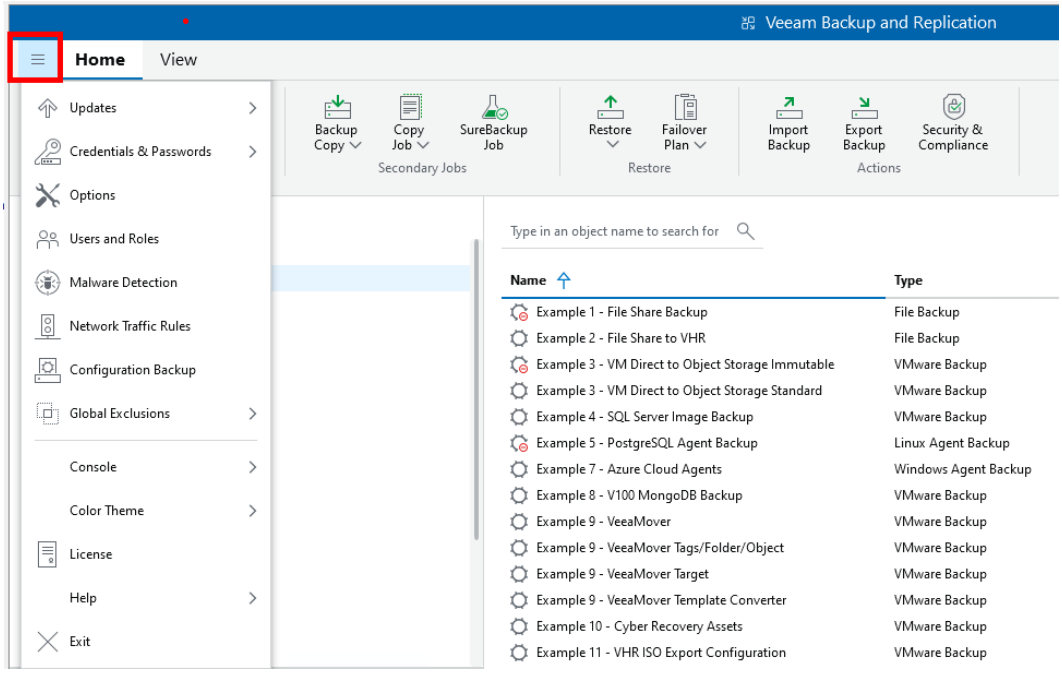


Figure 4.2: Main hamburger menu for options

What's rather helpful here is the number of different sections that provide visibility into what needs to be updated across the environment. The **Updates** area gives you a clear view of required updates (such as with components or missing updates), while the **Credentials & Passwords** section is an excellent way to rotate keys and manage credentials centrally. For a few versions now, it has been broken down into categories of credentials, making it easier to manage and know what goes where. The credentials and passwords are shown here:

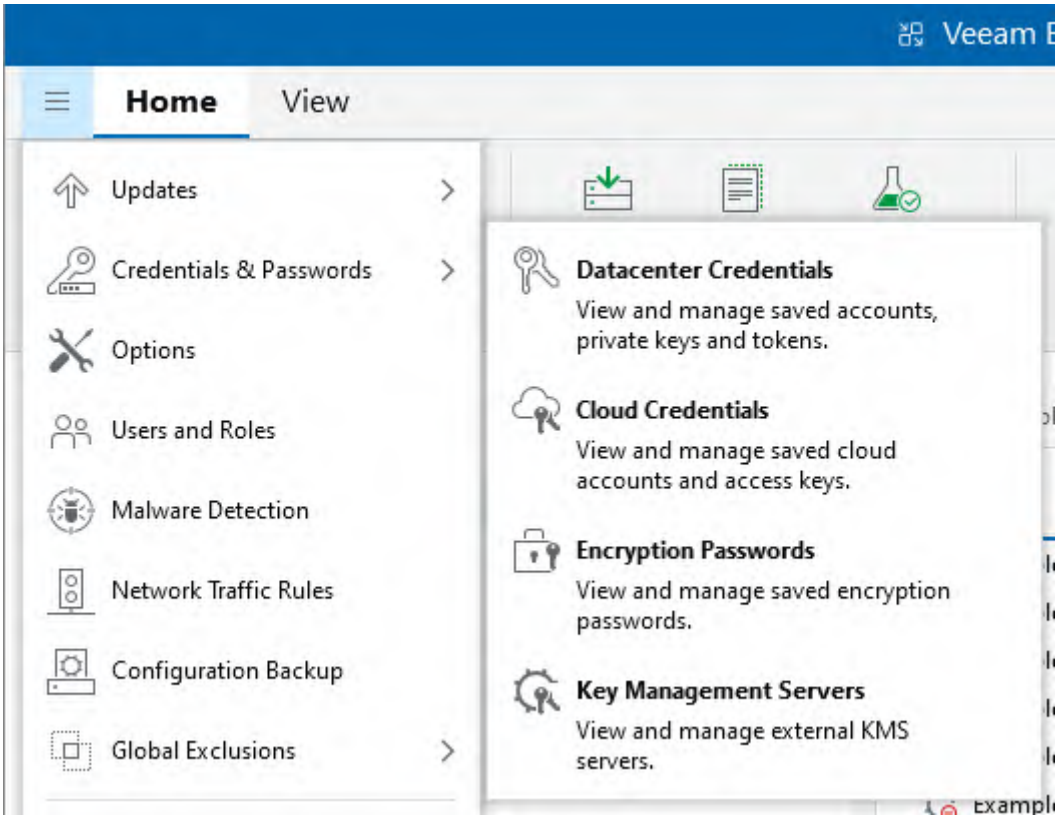


Figure 4.3: Credentials & passwords submenu options

Several global options can also be configured here. Along with **Users and Roles**, you can review **Malware Detection**, which I'll discuss in more detail later.

Somewhat surprisingly, my favorite area is the **Console** and **Color Theme** section of the main menu. Within the **Console** menu, you'll find shortcuts to several very powerful tools. The first allows you to launch a **Remote Desktop** connection to the Veeam Backup & Replication server itself. The second option opens a **PowerShell** console scoped directly to the server.

Technet24

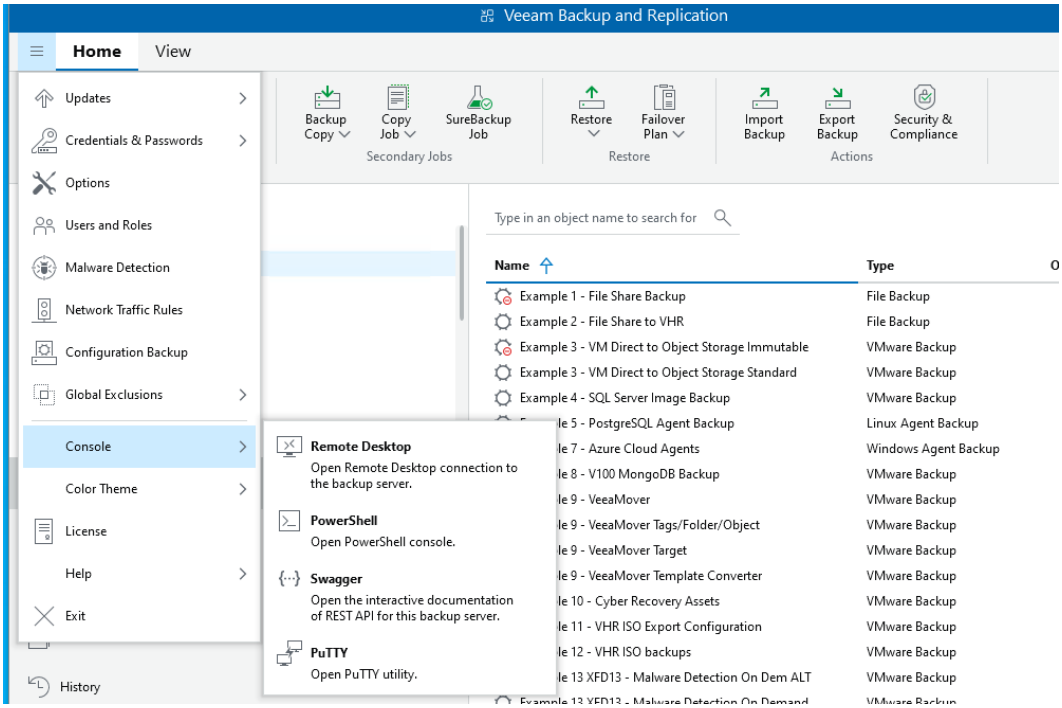


Figure 4.4: Console submenu options to select

You can also access the **Swagger**-based REST API documentation for the server. Finally, there's an option to launch **PuTTY** directly; if you don't already have the PuTTY SSH utility installed, you can still use it from here. This collection of tools is helpful for connecting to different systems, remoting into the Veeam Backup & Replication server, and running PowerShell commands without leaving the console.

The **Color Theme** option is purely cosmetic, but still extremely useful. It supports personal preferences as well as organizational standards, particularly for teams that prefer or require dark mode. In fact, dark mode is a long-standing customer request, and it finally arrived with version 13. For users who keep a persistent console open, such as in a network operations center, the dark theme option, labeled **Midnight Green**, can significantly reduce eye strain. Here is a view of Midnight Green:

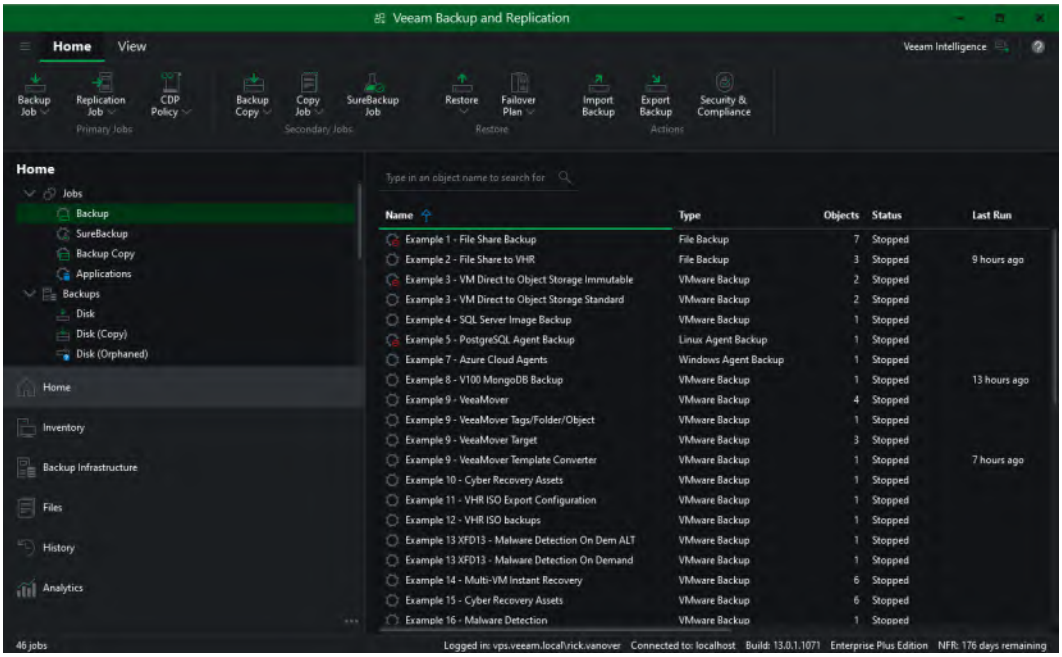


Figure 4.5: Dark mode version of the Windows console

You can also review **License** and **Help** information from the main menu. The license view, in particular, provides far more detail than the summary shown at the bottom of the console. From here, you can see how many days remain on the license, the license type, how many instances are in use, and you can even install a license if it's not managed by backup enterprise manager.

Technet24

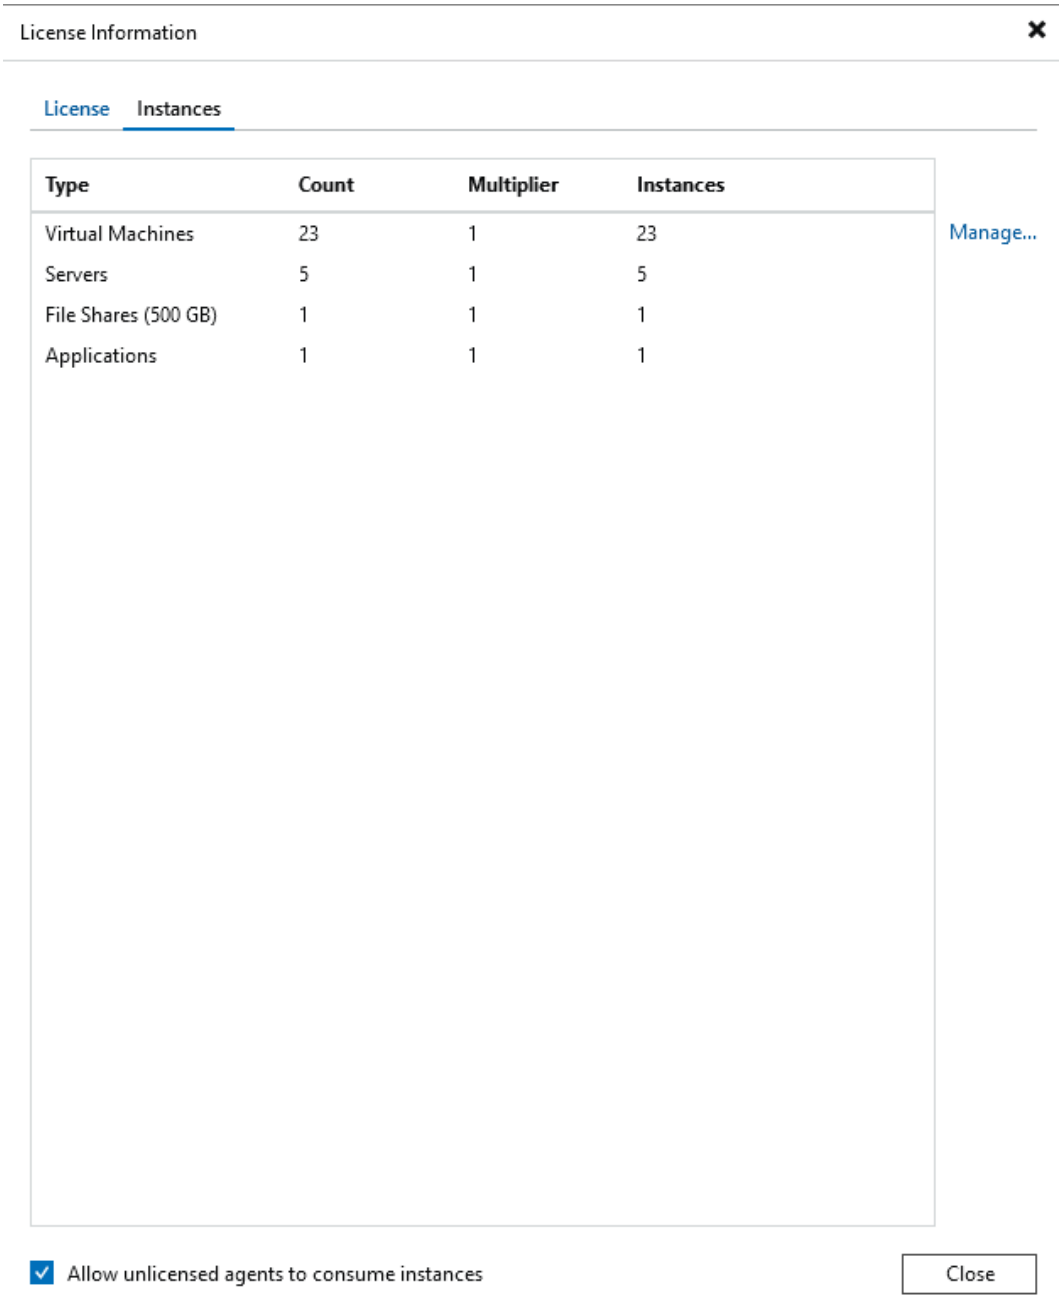


Figure 4.6: Veeam license information dialog showing instances used

Figure 4.6 shows what VUL is consumed on the **Instances** tab. You can see exactly which instance types have been consumed, along with their multipliers, making it easier to understand how license usage is distributed across the environment.

Jobs view

The **Jobs** view is where most administrators will spend much of their time. This view displays all job types within Veeam Backup & Replication and allows jobs to be sorted by type or viewed collectively. Because of this central role, it's important to apply thoughtful and consistent naming conventions.

In this example environment, there are many jobs, and the names are intentionally descriptive for demonstration purposes. In real-world environments, administrators often name jobs based on protection policies, target locations, applications, or business ownership. Regardless of the number of jobs in an environment, careful attention should be given to job nomenclature, so it is easy to understand what a job does, who owns it, and what it protects.

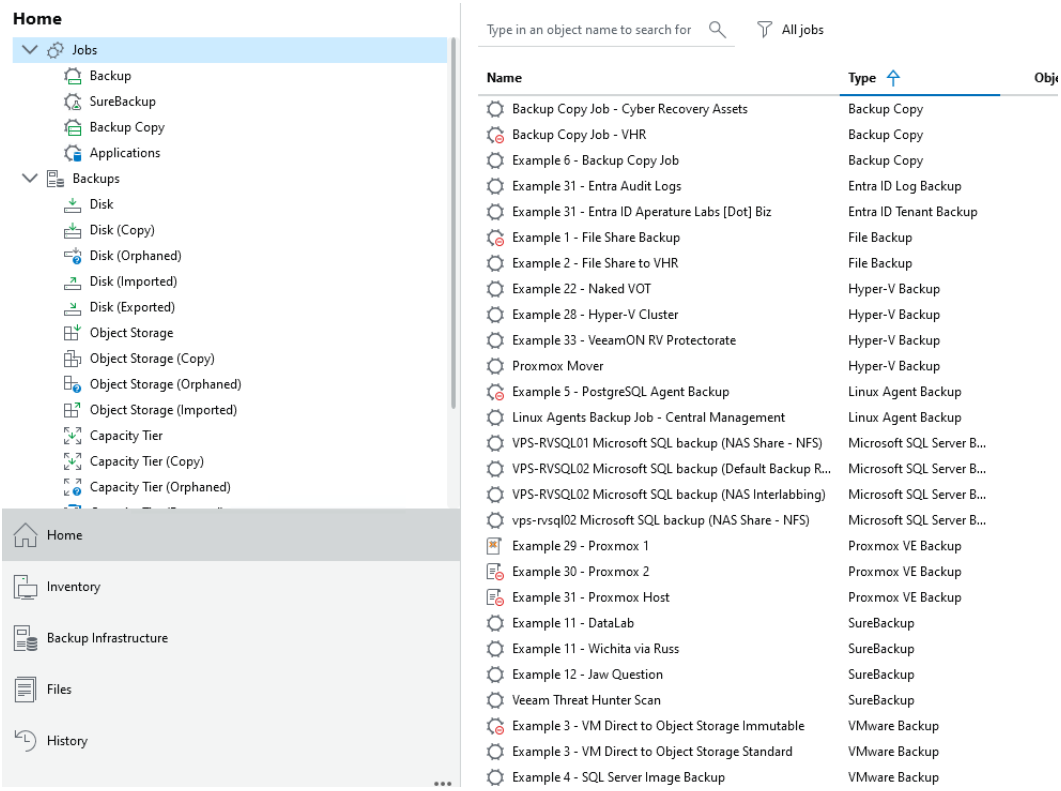


Figure 4.7: Jobs view for the Veeam Console sorted by Type column

Technet24

Figure 4.7 shows an example sorting by job type in the backup console. Each job includes a description field, which can be used to provide additional context without overloading the job name itself. Job names should remain short, easily identifiable, and consistent in structure, with more detailed information captured in the description. This approach helps reduce confusion, improves day-to-day usability, and simplifies ongoing administration at scale.

You can sort by different fields, such as name, last result, target (repository, replication resource, and so on) as well as the above example of job type. This often comes down to personal preference. Here is the same system, sorted by job name, which for my preference works well:

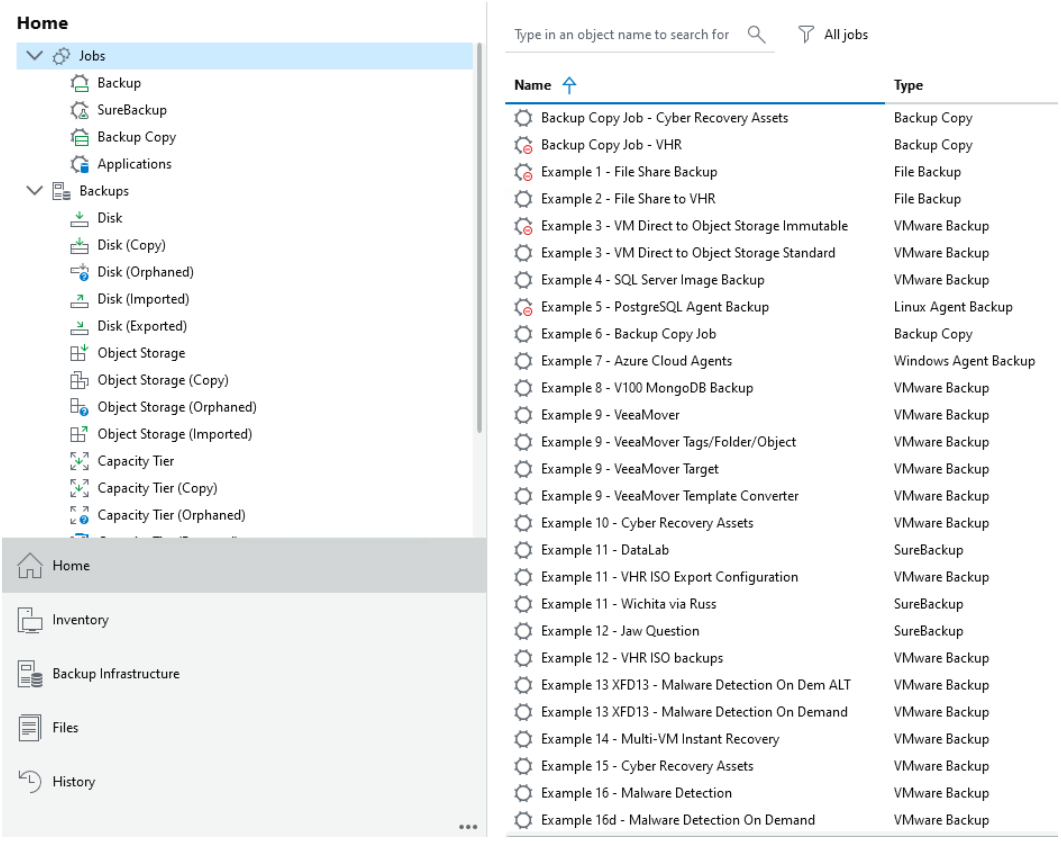


Figure 4.8: Jobs view sorted by Name column

Backups view

The **Backups** view (shown in *Figure 4.9*) is arguably one of the more powerful and underutilized areas of the Veeam Backup & Replication console. The following view is from a lab that I have access to:

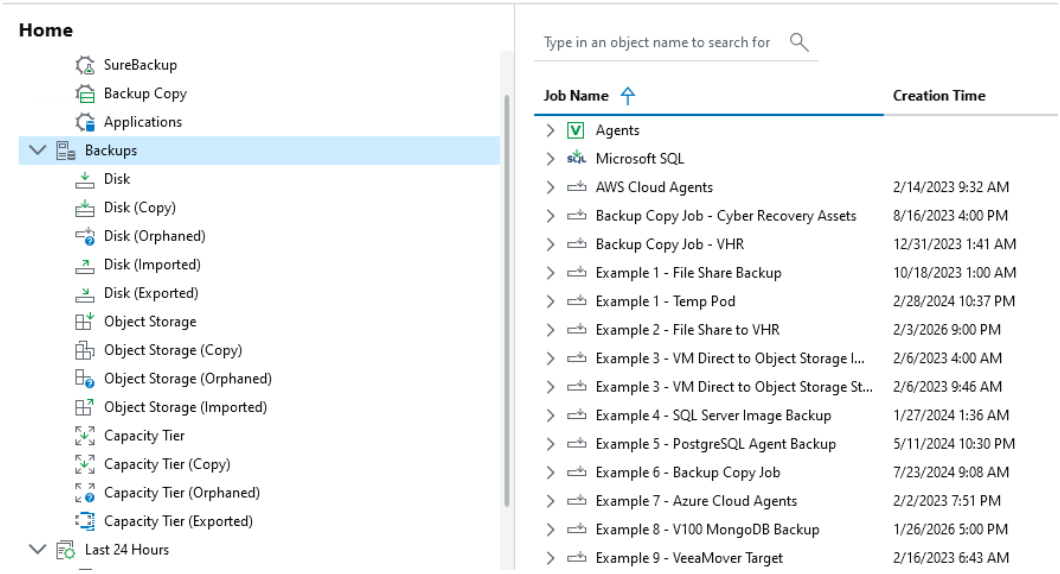


Figure 4.9: Backups view in the Veeam console

One important consideration when reviewing this view is the terminology used to describe different types of backups. **Copy** backups are created by backup copy jobs or through policies applied to a scale-out backup repository (SOBR). These are expected and typically part of a deliberate data placement or retention strategy. **Imported** and **Orphaned** backups, however, warrant closer attention. Imported backups result from bringing an existing backup file, which was possibly created on another Veeam backup server or even a different backup infrastructure, into the current Veeam Backup & Replication console. Orphaned backups indicate that the associated backup jobs have been deleted, while the backup files themselves remain on disk.

Note

Care should be taken before deleting imported or orphaned backups. Although they may initially appear to be candidates for cleanup, these backups might still be required, need to be relocated, or should remain intact for recovery or compliance reasons. Always inspect them carefully to understand their origin and purpose.

A powerful advantage is that these backups are fully usable. Even if the original job no longer exists, recovery operations can be performed directly from the **Backups** view as long as the files exist in the repository. *Figure 4.9* shows a MongoDB database in the **Backups** view with the recovery options identified:

Technet24

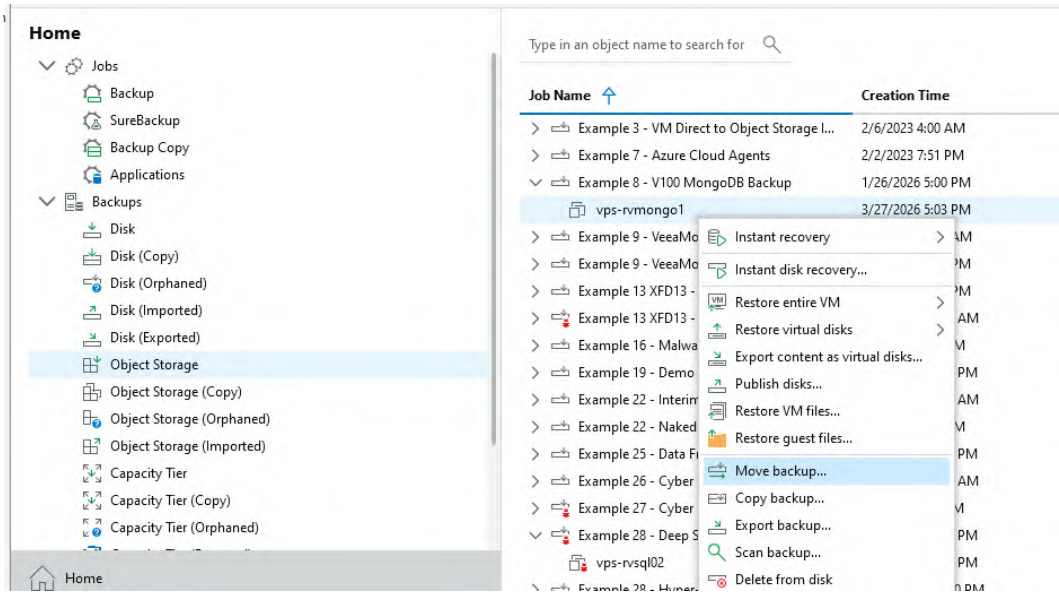


Figure 4.10: Recovery options for a particular backup

Last 24 Hours View

One of the most useful areas of the backup console user interface is the **Last 24 Hours** section. This view provides a quick summary of recent job activity, showing different job types, their session states, and overall status. It is especially valuable for quickly identifying jobs that failed, completed with warnings, or did not run as expected.

Figure 4.10 shows a view of the Last 24 Hours and a job being re-run that was not completed successfully. Here, the environment includes a few warnings that are acceptable, but it also highlights two job failures. This makes it easy to immediately see which jobs require further investigation. You can also observe that a replication job failed and is already running again within the same 24-hour window, indicating automatic retries or scheduled reruns.

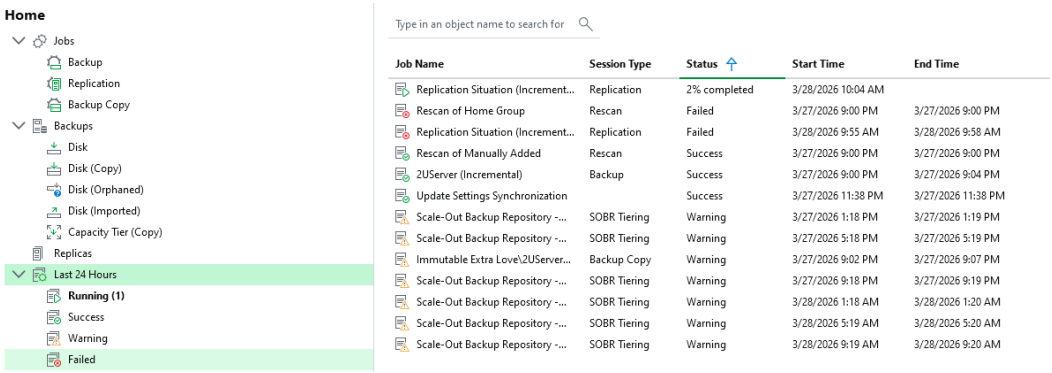


Figure 4.11: Last 24 hours view of jobs that have run

This view offers a concise yet information-rich snapshot of how jobs have behaved over the past day, including their session types and outcomes. Reviewing the **Last 24 Hours** section daily is a recommended best practice, complementing email alerts by providing visual context and enabling faster operational awareness directly within the console.

Now we will take a look at the inventory view of the Veeam console and what is protected by Veeam.

Inventory view

The **Inventory** view is an excellent way to see a comprehensive snapshot of what is protected by Veeam. This includes virtual platforms, physical agents, cloud workloads, unstructured data, and more—all presented in a single, unified view. It allows administrators to quickly understand coverage across the entire backup environment. This view is also where **Malware Detection** results are surfaced, making it an effective place to identify potential security concerns across protected data. If any data is flagged as suspicious or confirmed to contain a threat, it is clearly visible within this interface.

Figure 4.11 shows the view of an environment with a lot of malware detection events as well as several different platforms inventoried for backup:

Technet24

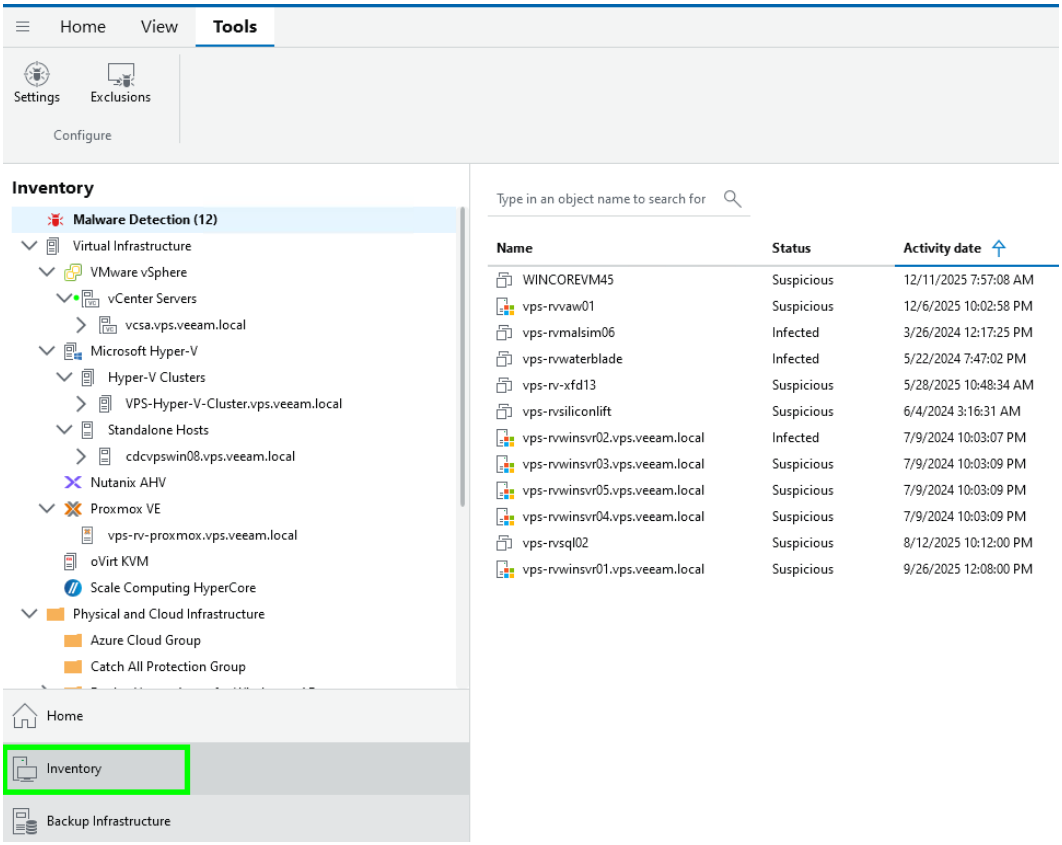


Figure 4.12: Inventory view of the Veeam Console

In this example, there are twelve malware detections. Many of these are simulated to demonstrate capability, but they illustrate how malware detection can quickly highlight areas that may require further investigation. In real-world scenarios, this functionality provides valuable insight into whether Veeam has detected potentially malicious activity within protected backups, helping organizations respond more quickly and confidently.

As you can see, the inventory view allows you to check malware detections and create protection groups. Now let's look at the backup infrastructure view where you will see backup proxies and backup repositories.

Backup infrastructure view

The **Backup Infrastructure** view is a critical component of day-to-day administration in Veeam Backup & Replication. This is where core infrastructure components such as backup proxies and repositories are configured and managed. It is also where scale-out backup

repositories (SOBRs) are defined, along with associated policies for capacity tiering, immutability, and storage-specific behavior.

Because this view serves as the foundation of the backup environment, consistent naming conventions and disciplined administrative practices are especially important. Clear nomenclature helps ensure maintainability and reduces confusion as the environment grows.

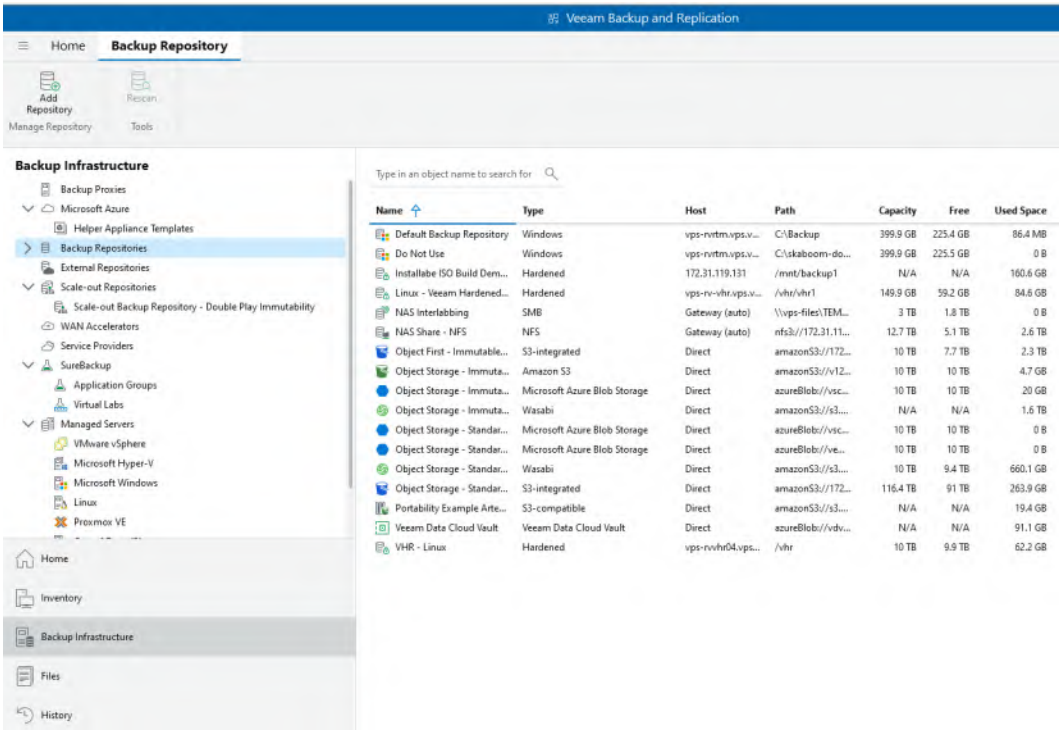


Figure 4.13: Backup Infrastructure view of the Veeam console

The **Backup Infrastructure** view is also where connections to service providers are established and where components required for **SureBackup** jobs are configured. SureBackup plays a key role in automated recovery verification, ensuring that backups are not only successful but also recoverable.

Not every environment will use all available infrastructure components, but it's important to understand that when specific roles or services are required to enable certain Veeam features, this is where they are configured. As a result, the **Backup Infrastructure** view becomes the central location for enabling, extending, and validating the capabilities of the overall backup platform.

Technet24

Files view

The **Files** view in the user interface is one of the more well-known and historically significant parts of the product. This is where **FastSCP** originated and where its capabilities continue to exist today. FastSCP was one of the technologies that helped put Veeam on the map in the early 2000s. It solved a common and frustrating problem at the time: moving files such as ISO images from Windows systems into the closed VMware VMFS file system, which was not easily accessible through standard tools. FastSCP enabled simple, efficient file transfers between systems, addressing a real operational challenge for administrators working with early VMware environments. Those capabilities have since been fully incorporated into the **Files** view within Veeam Backup & Replication.

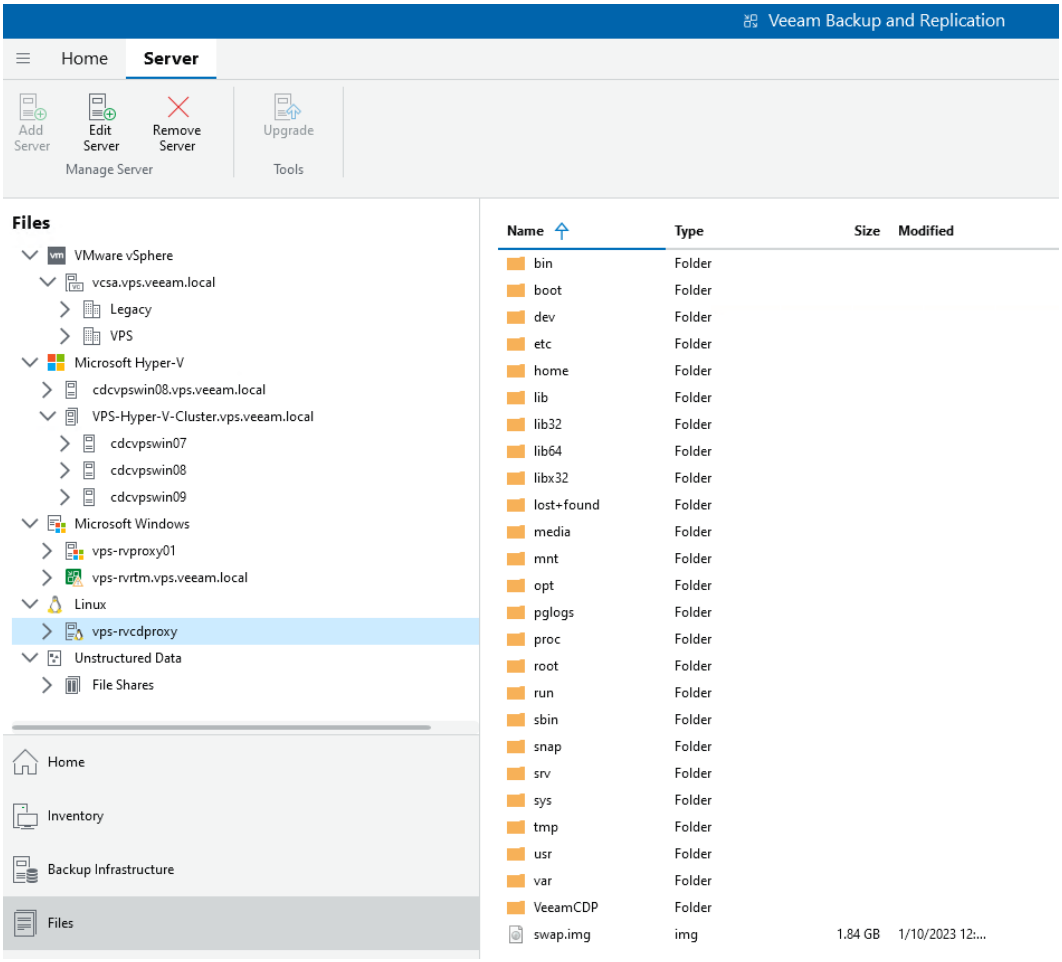


Figure 4.14: Files view of the Veeam console

Today, the **Files** view still provides powerful file portability across platforms. For example, you can copy files from a local Windows path directly into a VMware datastore or transfer files from a Linux host into a Hyper-V cluster path. This flexibility makes it a versatile tool for moving data between otherwise disconnected environments.

In conclusion, the Files view is like Windows explorer for navigating the file system of servers and shares. It is a handy way to do any file copy operations across managed systems in the Veeam environment. Now let's look at the process of updating the Veeam console when there is a patch installed on the main Veeam Backup & Replication server.

Updating the Veeam Backup & Replication user interface

Like all software, Veeam Backup & Replication requires periodic updates. Over time, you will need to update both the backup server components and the Backup & Replication console application.

> **Note**
>
> It's important to note that the backup server and the console must be running the same version. Keeping all components aligned and up to date ensures compatibility, stability, and access to the latest features and fixes.

In environments that use the Veeam software appliance, updates are handled automatically. For those familiar with Veeam, this is new behavior since version 13 has been released. These updates are mandatory and are installed without manual intervention. After an appliance update occurs, the next time you connect using the Veeam Backup & Replication console, you will be prompted to update the console to the matching version. This console update process is typically straightforward and completes quickly.

Maintaining version consistency across components is a simple but critical practice that helps avoid unexpected issues and ensures a smooth administrative experience. It is reasonable to expect that on one random day, an update will be installed, and due to that the console may need to be updated as well. If you are using VSA, updates will happen automatically and you'll get a message like this occasionally to do an update of the console. This can take a few minutes but is otherwise straightforward:

Technet24

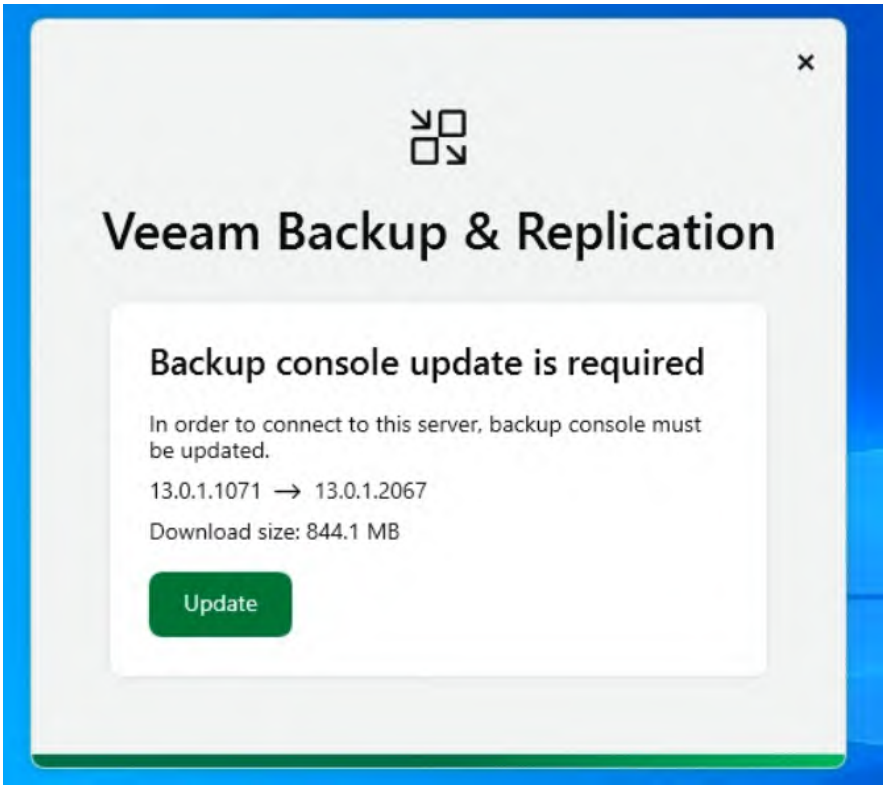


Figure 4.15: Veeam console update prompt and update to proceed

Overall usage of the Windows console is part of the story since version 13 has been released with the new web user interface. It's important to note for a few generations of the product that both will be needed, so it is important to be familiar with both.

Web UI vs. Windows UI in v13

The web user interface is a strong option for accessing Veeam from a variety of devices, including tablets and systems running non-Windows operating systems. For the foreseeable future, however, usage largely comes down to personal preference and the task at hand. While I personally find the web UI more pleasant to use, it does not yet offer full feature parity with the Windows UI. As a result, I primarily use the web UI to monitor the environment and check infrastructure status, while relying on the Windows UI for configuration tasks. In fact, on first logon you'll be given this message:

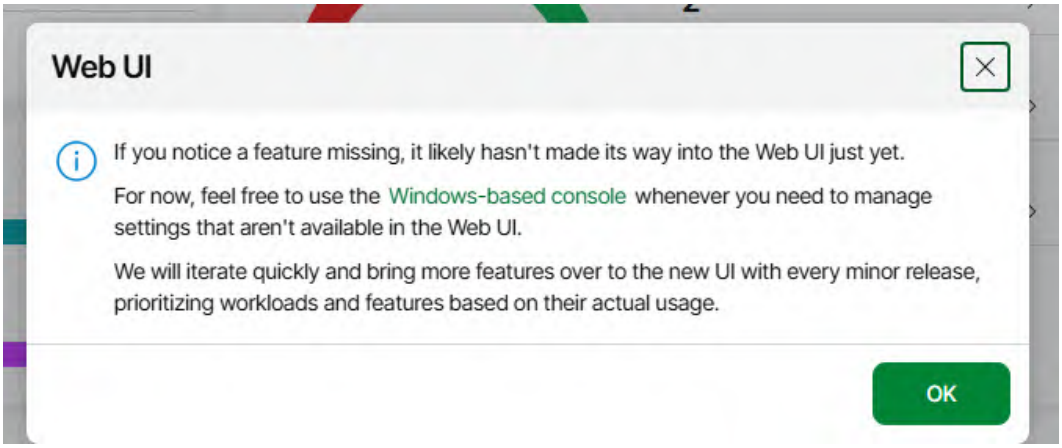


Figure 4.16: Web UI notice about feature parity to the Windows console

Common operations such as creating backup and replication jobs can be performed on both interfaces. However, recovery scenarios currently favor Windows UI, which provides the most complete set of restore capabilities. All restore workflows in the Windows UI are fully implemented and leverage familiar tools such as Windows Explorer to guide users through recovery options.

It's important to note that restore functionality in the web UI will continue to expand over time. Users familiar with the Windows UI will also find the web UI intuitive, as the overall workflows and concepts align closely. Directionally, there is a significant investment in enhancing the web UI and moving it toward feature parity with the Windows UI. While no specific timeline has been disclosed, it is reasonable to expect improvements with each major release. At the same time, the Windows UI remains fully supported and is not being deprecated. *Figure 4.16* shows a view of the web user interface in the 13.01 build:

Technet24

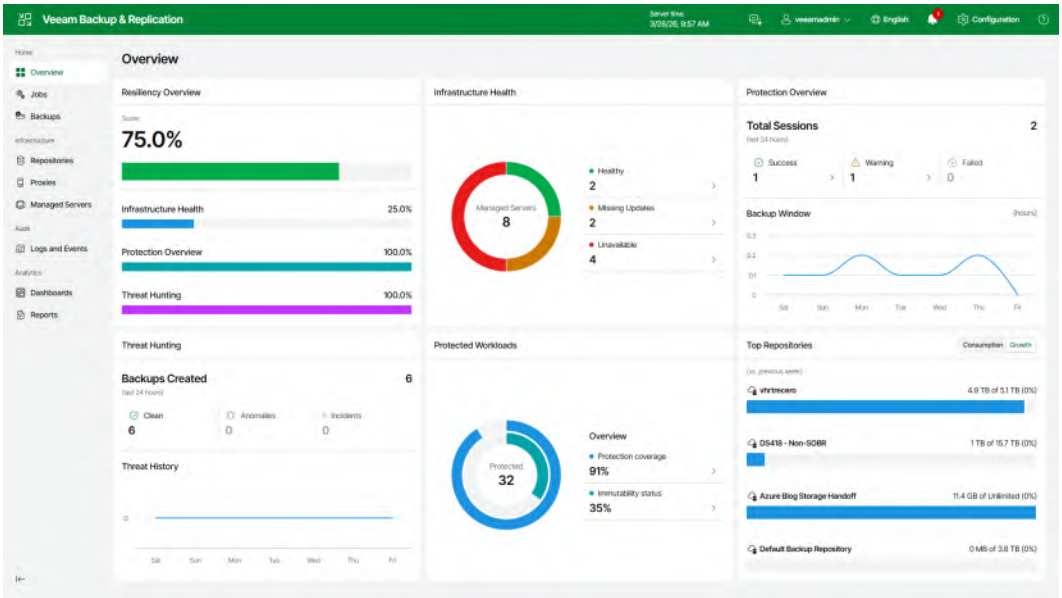


Figure 4.17: Overview screen of the web UI

One way to interpret the differences between the web UI and the Windows-based backup console is that the web UI provides a very good at-a-glance look at the status of the backup infrastructure. And the Windows-based backup console is meant to go and do tasks such as creating a backup job, performing a recovery, and so on. When looking at the same server in the Windows-based backup console, the view is different as shown in the following figure:

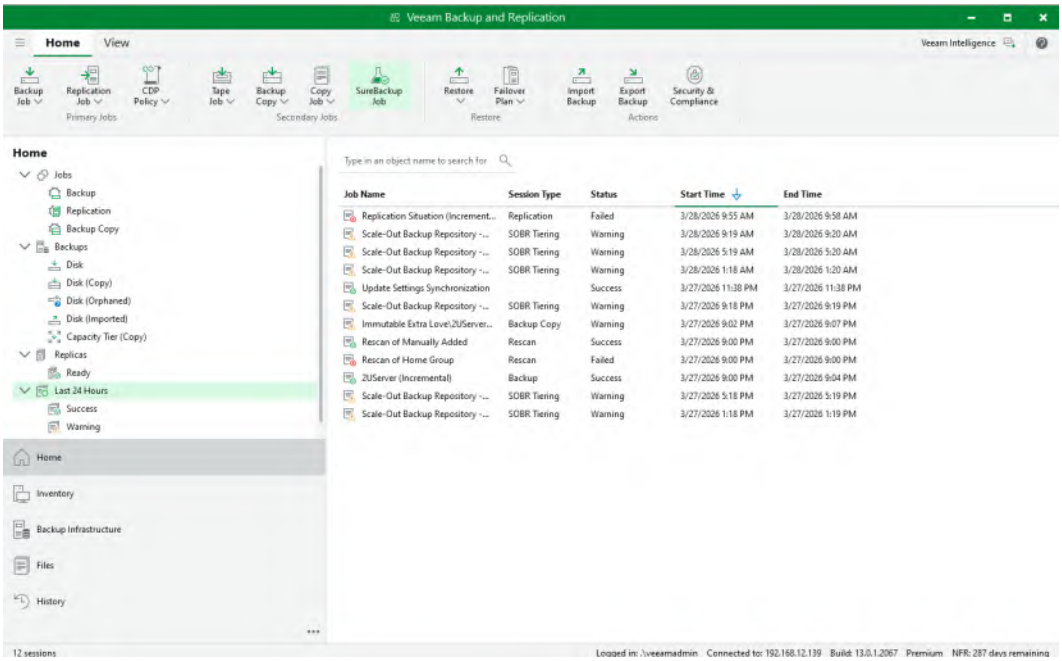


Figure 4.18: Veeam console comparison view to Web UI

We have now had a good look at the differences between the Web UI and Windows UI in Veeam Backup & Replication 13. Now we will explore some advanced tips and tricks you can use while working in the Veeam Console on Windows.

Advanced tips and tricks for using Veeam Console

One of my favorite aspects of working at Veeam is the sheer expertise that surrounds me. I am particularly in awe of Veeam tech support, so much so that I have even hired two folks to my team based on their skill and demeanor that was founded by their time in Veeam support (Shout out to *Karinne Bessette* and *Andrey Stadler*). In preparing this chapter, I had the opportunity to poll some tech support leads for some tips and tricks to share and I want to give a special thanks to: *Elena Soboleva*, *Matt Nelson*, *Paul Loewenkamp*, *Spencer West*, *Tim J. Smith*, and *Ken Sauer*.

Best practices for backup consoles

One of the most important best practices when using the Veeam Backup & Replication console is to close it when it's not in use and explicitly log out of active sessions. From a security perspective, organizations should also enable multi-factor authentication and four-eyes authorization to provide additional layers of protection. It's natural to spend time exploring

Technet24

the user interface to learn more about the product, especially when preparing for less common tasks such as specialized recovery workflows or advanced features. This curiosity is a good thing and often leads to more effective use of the platform.

Administrators may occasionally encounter warnings related to display resolution or DPI scaling. Personally, I prefer running at 100% DPI scaling, but this isn't always practical in environments with lower screen resolutions. These challenges are especially common when using jump boxes, where frequent disconnects and reconnects can cause display resizing issues. This is another reason why logging off and closing the console rather than leaving sessions open is strongly recommended.

Beyond tooling and security, organizational readiness is equally important. Every environment should have at least one or two additional team members trained to use the Veeam console. Incidents rarely occur at convenient times, and recovery may be required when primary administrators are unavailable. Ensuring others can confidently perform tasks such as running backups, executing recoveries, or even conducting advanced operations like threat hunting scans is critical.

A particularly effective practice is to intentionally test recovery scenarios when key individuals are unavailable. This helps validate cross-training, exposes gaps in procedures, and ensures the broader team is prepared to respond effectively when it truly matters.

Session statistics

One of my favorite tips and tricks to share from the user interface is session status history with the left and right arrow keys (← and →) when viewing the statistics of a job. Simply click a job, then click statistics to use this nice feature. For example, if you are troubleshooting a job and you want to see the history, simply hit the ← (left arrow) to see the previous job and the → (right arrow) to see the next job. In *Figure 4.18*, I'm running my current job after a change to diagnose why the replication job failed:

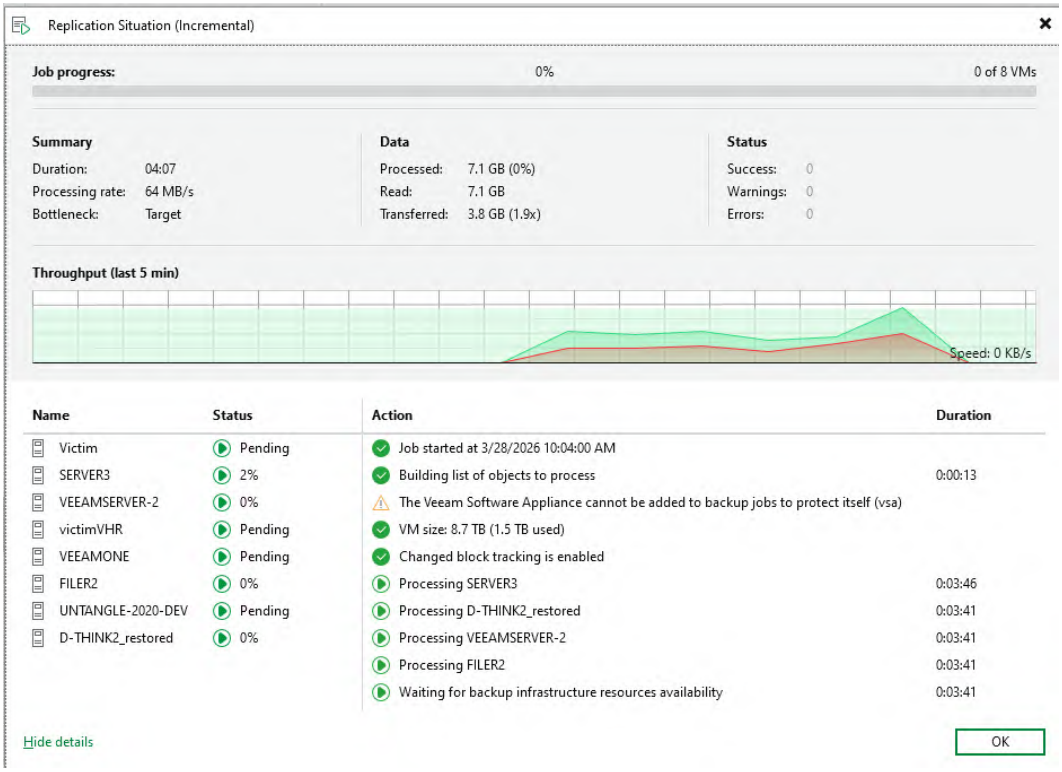


Figure 4.19: Backup job statistics dialog

When I hit the left arrow, it'll take me back to the previous session and I can see that it failed. One trick is to look at the top left to see the "**Job started at**..." field to orient yourself to the timestamps. In the following figure, you will see it is earlier than the 10:04:00 AM start which was the job that failed, and that my change has corrected the issue:

Technet24

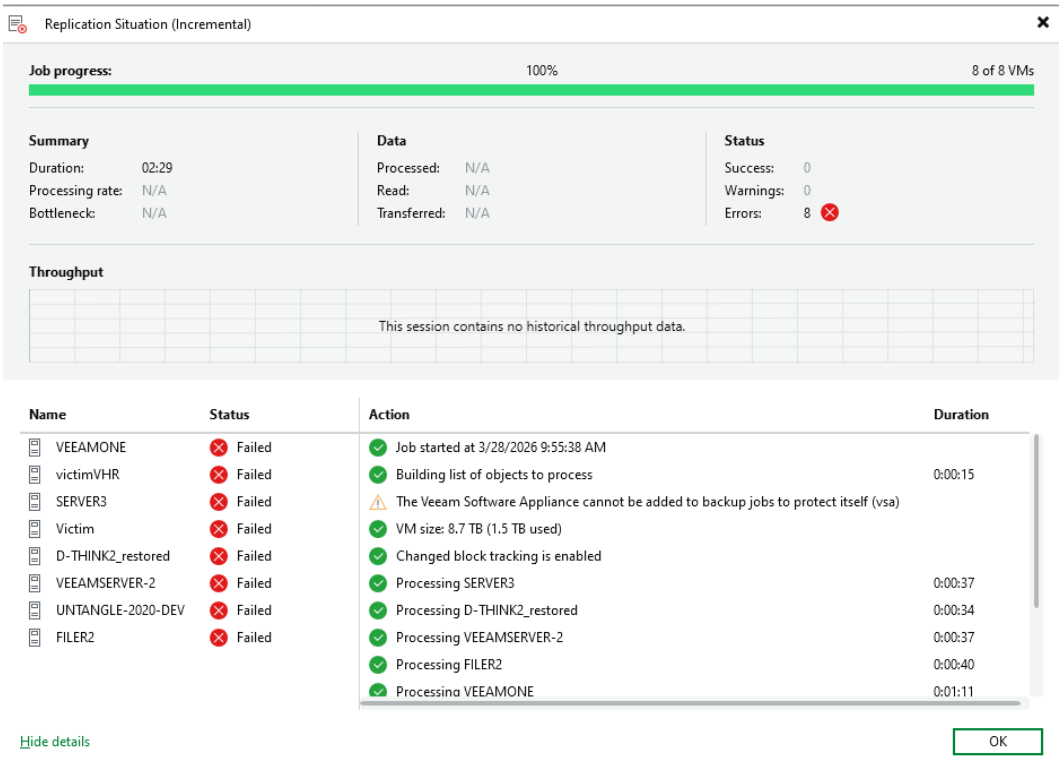


Figure 4.20: Job statistics window with successful backup

What is best about this left and right arrow keys trick is that you can do it while jobs are running as well as when they are not. It all starts by looking at the statistics of a job. If anyone wonders, the fix above did the trick and the replication job was completed successfully!

Text editing, FAST

The original text editor that was part of FastSCP is still in the Veeam Backup & Replication console, and in many cases, it is much easier than remotely connecting to a system (Windows or Linux) to do a quick view or edit. You can right-click on a file and do an edit directly:

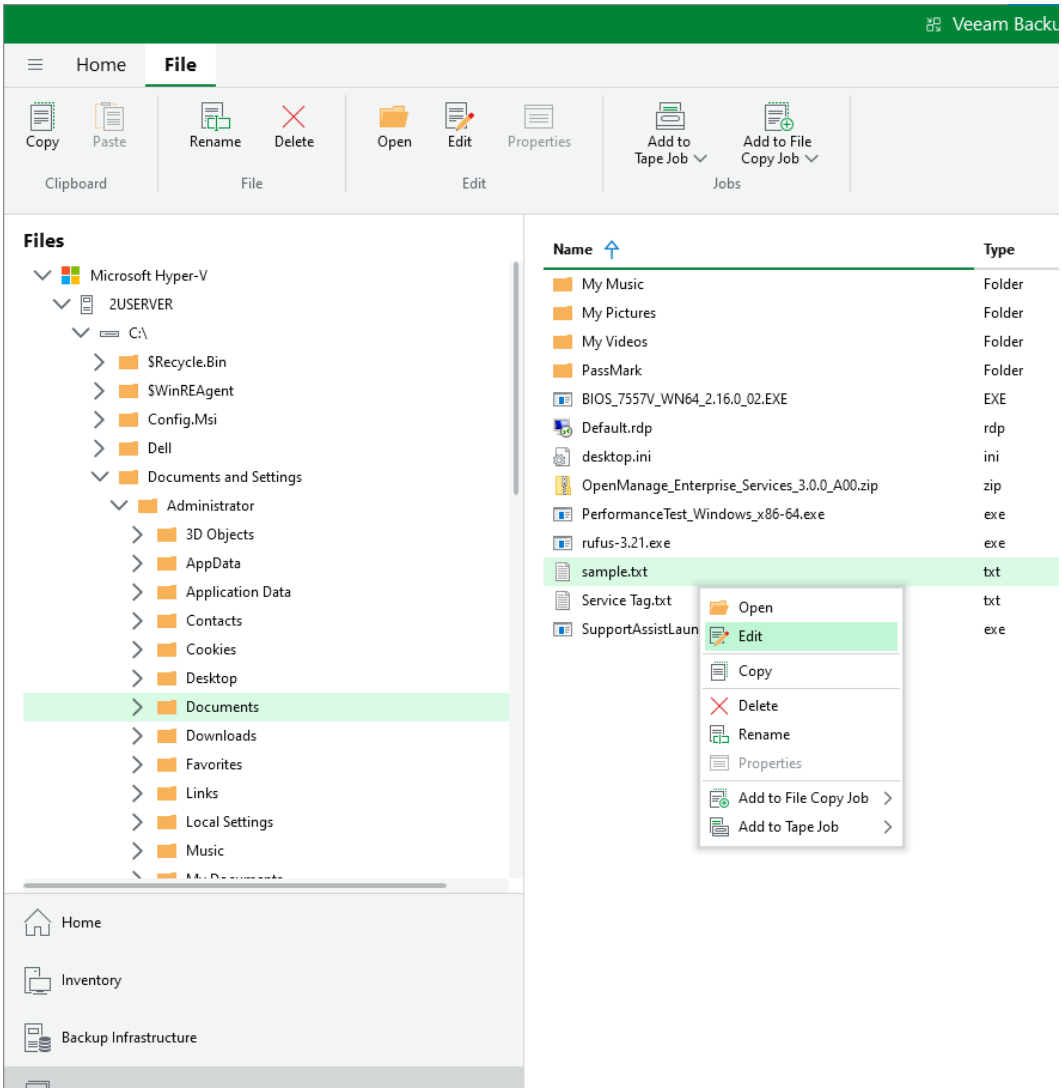


Figure 4.21: Text file editing within the Veeam console – Files tab

The editor even has some find and replace functionality as well as the ability to save changes. Note that edit capability is prohibited if four-eyes authorization is in use. The editor is shown in the following figure:

Technet24

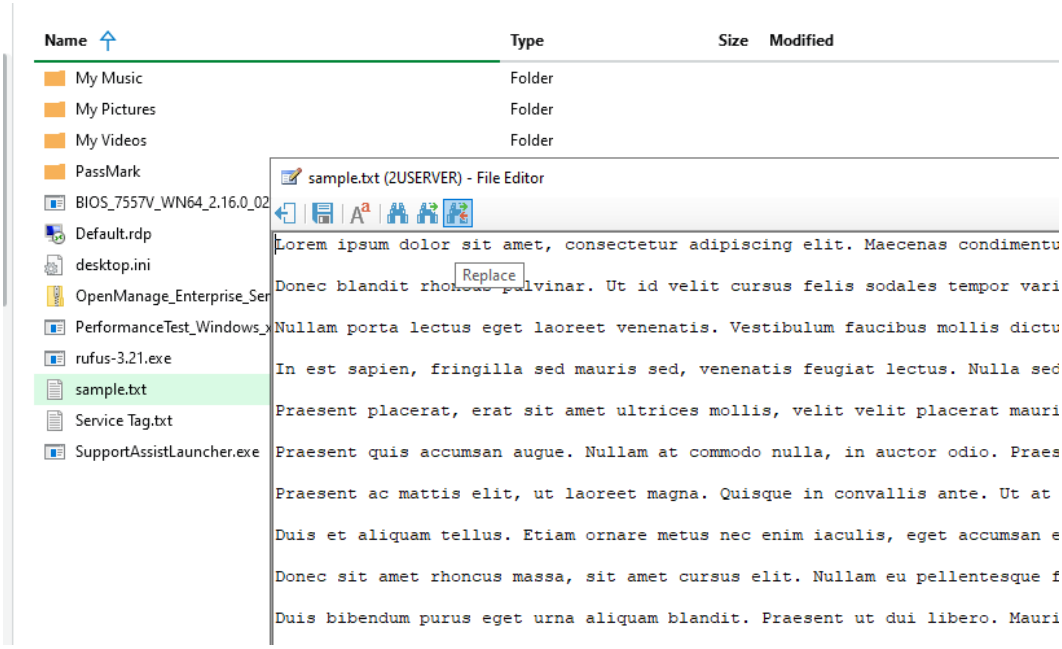


Figure 4.22: Editing a text file from within the Veeam console in the file editor

Secret CDP bottleneck information

All jobs have bottleneck analysis views, but there are many more for CDP Policies. If you are using CDP, there are additional infrastructure bottlenecks that are reported for policies. If you hover over the bottleneck on a CDP session, you will see six different categories reported:

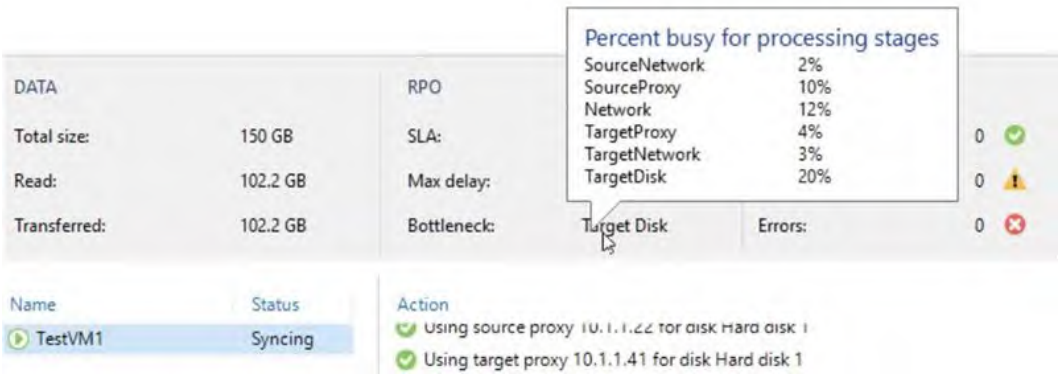


Figure 4.23: Bottleneck statistics within a job

Backups information

The backups view node will show you all the backups across all the types/storage locations, whereas each individual node underneath filters things out. What you see under **Backups | disk** and **Backups | Capacity tier** are different, but what you see under **Backups** is both 2. Same goes for replicas. If you navigate to just **Replicas** and not **Replicas | Ready** you can see all existing replicas even if they are in faulty state. This filtering applies to what you see inside the backup chain properties once you open it. Here is one where you can view details including how long a backup chain file is immutable until:

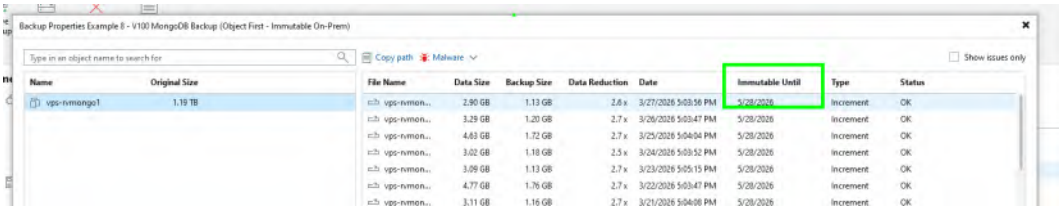


Figure 4.24: Backup properties showing restore points

Hello CTRL

At many points in the Windows UI, some commands only appear if you press the CTRL key before a click or right-click. One example is the **Remove from** configuration option as shown for this orphaned backup, which I can remove from the database:

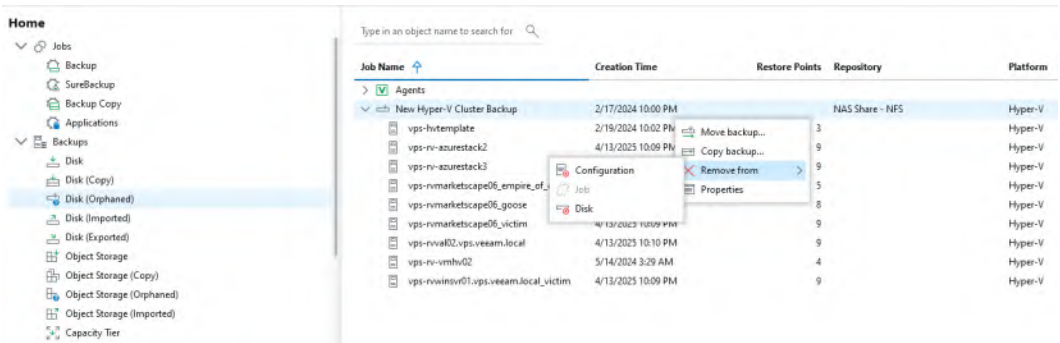


Figure 4.25: Using the CTRL key to unlock more options

Tip

I need help!

On most parts of the Veeam Backup & Replication console, you can press the F1 key to go to the relevant page in the user guide for that part of the product. The link to the

Technet24

help center is the following: `https://helpcenter.veeam.com`, but is handy to go right to the specific area!

Refresh first, avoid questions later

When you add virtual machines to an existing job, it is recommended to use the **Refresh** button first, so the backup wizard will forcibly update infrastructure details to bring the latest inventory in.

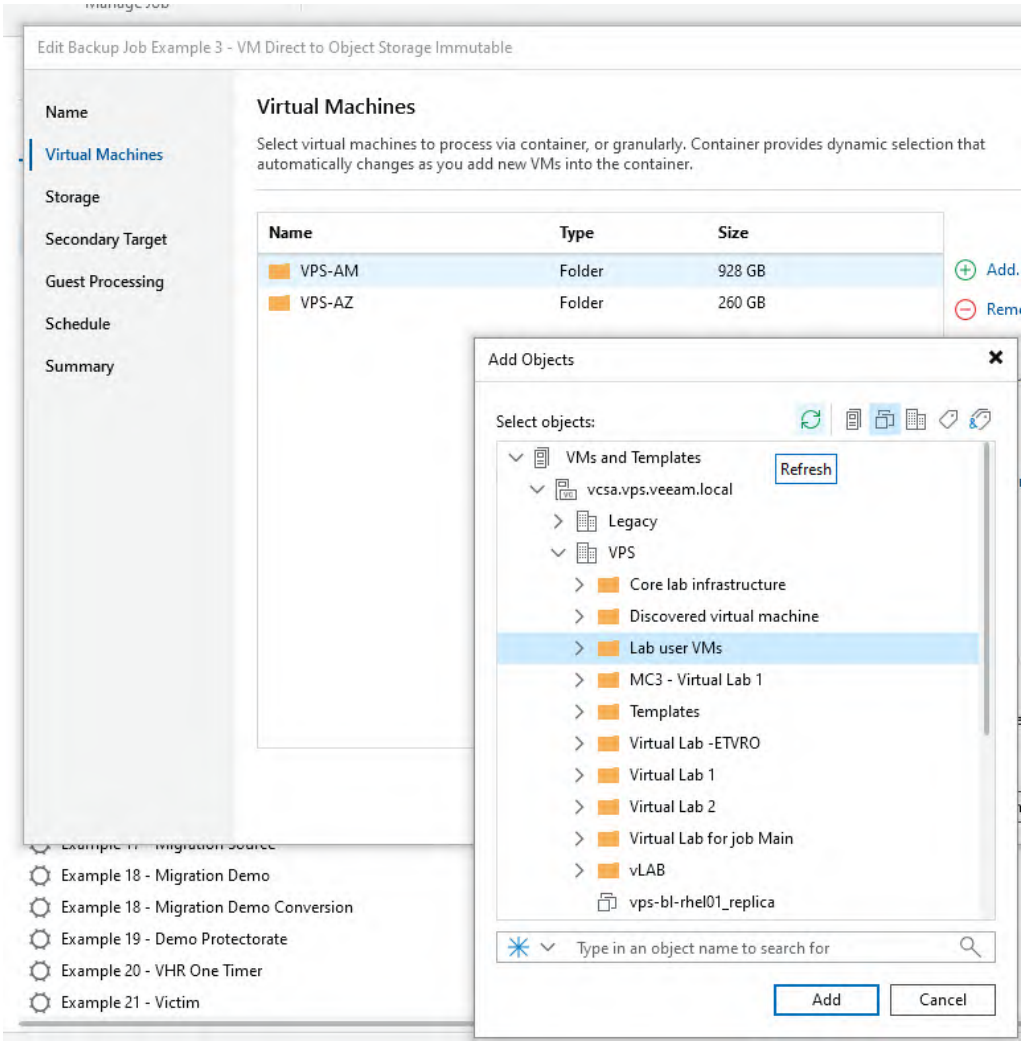


Figure 4.26: Editing a job to add more for backup

The Veeam Backup & Replication console will remain the primary day-to-day administration platform in version 13 for the foreseeable future. While the web UI continues to see meaningful improvements, it is reasonable to expect that the full set of administrative capabilities will continue to reside in the Windows-based Veeam console. This is especially true when you consider many of the advanced tips, workflows, and troubleshooting techniques that administrators rely on today, which will likely take some time to be fully incorporated into the web UI.

That said, the Windows console is highly customizable for daily operations. Administrators can tailor views to focus on the information that matters most to them, making it well suited for both routine management and deeper operational tasks. It's also important to note that the console must be kept up to date as Veeam Backup & Replication continues to evolve. As new updates and enhancements are delivered, corresponding updates to the Windows console will be required. This update cadence is ultimately a positive development, particularly from a security perspective. Keeping the console current helps maintain a reduced attack surface, ensures compatibility with the latest platform improvements, and allows new capabilities to be rolled out more quickly and consistently.

"*Veeam is too hard to use*," said nobody ever, and the fact that most workflows in Veeam Backup & Replication are seven (or fewer) clicks of a wizard speaks to that truth.

Summary

Veeam Backup & Replication console is where a Veeam user will spend a lot of time. This chapter examined the backup console in version 13 and its role as the primary interface for administering and operating the Veeam data protection platform. We began with a brief look at the evolution of the Veeam user interface, highlighting how early design decisions and usability principles have resulted in a consistent, wizard-driven experience that continues to scale with modern infrastructure requirements.

We then covered how to deploy and access the console, including installation options for Windows-based environments and VSA, along with considerations for remote administration. From there, we explored console navigation and core functional views, such as home, jobs, backups, inventory, backup infrastructure, and files, showing how each area supports everyday administration, monitoring, and troubleshooting activities. The chapter also addressed console updates and version alignment, emphasizing the need to keep the console synchronized with the Veeam Backup & Replication server to ensure stability, security, and access to new features. A comparison of the Windows UI and web UI in v13 clarified how the two interfaces complement each other today, with the Windows UI remaining essential for advanced configuration and recovery tasks while the web UI continues to evolve.

Technet24

Finally, we reviewed advanced tips and best practices to improve efficiency, security, and day-to-day usability, providing practical guidance for administrators managing production environments. Together, these topics prepare you to confidently use the Veeam Console as a central management tool in Veeam Backup and Replication v13.

With a solid understanding of how to operate and navigate the Veeam Console, the next chapter introduces Veeam AI – Veeam intelligence, exploring how AI-driven insights enhance visibility and decision-making across the Veeam Backup & Replication environment.

Further reading

Veeam Backup & Replication Console - `https://helpcenter.veeam.com/docs/vbr/userguide/backup_console.html?ver=13`

Part 2

Veeam AI, Best Practices for Active Directory, High Availability (HA) & Veeam CDP Enhancements

The objective of this section is to dive into the new Veeam AI – Veeam Intelligence that is built-in to the program, allowing you to get troubleshooting assistance and many other things. This section will also include a look at how to best use Active Directory when working with a Veeam infrastructure. We will take a deep drive into the new High Availability (HA) feature within the new version. We will also discuss the Veeam CDP enhancements in Veeam Backup & Replication 13.

This part of the book includes the following chapters:

- *Chapter 5, Veeam AI – Veeam Intelligence*
- *Chapter 6, Best Practices for Active Directory Domains with Veeam*
- *Chapter 7, Veeam High Availability (HA)*
- *Chapter 8, Veeam CDP – Universal Continuous Data Protection*

Technet24

5

Veeam AI – Veeam Intelligence

Veeam Intelligence is part of Veeam Backup and Replication v13, the latest release of Veeam, and allows you to interact with a "ChatGPT" type interface that Veeam manages. This chapter examines what Veeam Intelligence is, what is new, and how you, as a backup administrator, can leverage it to help with your daily tasks. We will also examine how you access it through WebUI or Windows console. We then look at the configuration modes in which Veeam Intelligence can operate. Lastly, we will examine how Veeam Intelligence makes the life of a backup administrator easier.

In this chapter, we're going to cover the following main topics:

- Veeam Intelligence – A new way to help
- Accessing Veeam Intelligence
- Veeam Intelligence operation modes
- Veeam Intelligence: AI-powered assistance for backup administrators

Technical requirements

To ensure you can follow along with this chapter, you must have access to the WebUI or the Windows console.

You can find the Veeam Intelligence button in the Veeam console as shown in *Figure 5.1*.



Figure 5.1: Veeam Intelligence in the console

Technet24

Accessing Veeam Intelligence from the Veeam console is very straightforward via the toolbar button. You can access it through WebUI by clicking on the Veeam Intelligence button in the toolbar.

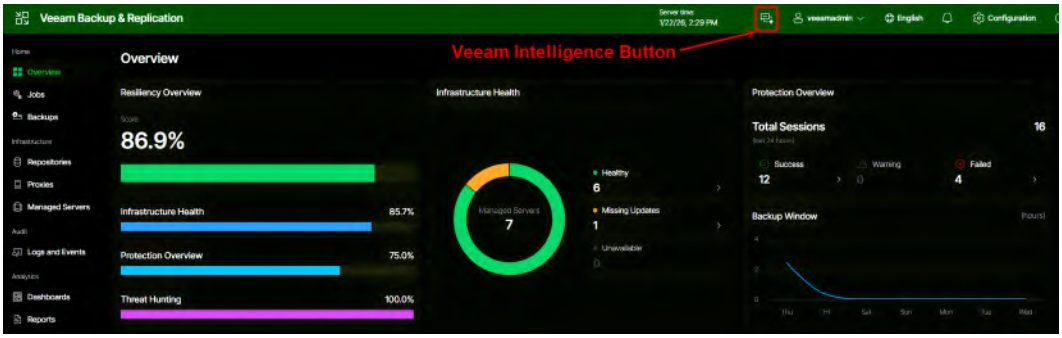


Figure 5.2: Veeam Intelligence in the Web UI

Veeam Intelligence – A new way to help

Veeam Intelligence represents a fundamental shift in how organizations interact with their data protection infrastructure. As the Gen AI core of the Veeam Data Platform, Veeam Intelligence transforms raw telemetry and complex backup data into clear, actionable insights delivered through natural language conversations. Rather than requiring administrators to navigate complex dashboards, query databases, or interpret cryptic log files, Veeam Intelligence enables users to simply ask questions about their environment and receive immediate, contextual responses.

This capability addresses a critical challenge in modern data protection: the overwhelming volume of information generated by backup operations across hybrid and multi-cloud environments. As organizations manage increasingly complex infrastructures spanning physical servers, virtual machines, containers, cloud workloads, and SaaS applications, the ability to quickly assess protection status, identify risks, and make informed decisions has become essential to maintaining data resilience.

Understanding Veeam Intelligence

Veeam Intelligence is a built-in conversational AI assistant integrated directly into the Veeam Data Platform and Veeam Data Cloud. Unlike generic large language models trained on public internet data, Veeam Intelligence is purpose-built to understand your specific backup environment. It connects directly to your Veeam infrastructure to provide contextualized answers based on your actual job configurations, protection policies, storage repositories, and historical performance data.

The system operates through multi-agent architecture, with specialized AI agents handling different types of queries and analysis tasks. This design allows Veeam Intelligence to provide both conversational guidance and deep technical analysis while maintaining appropriate security boundaries and access controls. Every action is logged with a detailed audit trail, and access to backup data is scope-limited to ensure Veeam Intelligence operates only within approved boundaries and datasets.

At its foundation, Veeam Intelligence leverages Azure OpenAI services in partnership with Microsoft to deliver advanced AI-driven guidance. This partnership ensures enterprise-grade security, compliance, and reliability while providing the sophisticated natural language processing capabilities required to translate complex technical concepts into human-readable insights.

What's new in version 13

Veeam Data Platform v13 introduces significant enhancements to Veeam Intelligence, expanding its capabilities through new AI agents, improved visualization features, and deeper integration with threat detection systems. The release introduces advanced mode, which unlocks access to specialized AI agents designed for specific analytical tasks. We will now look at these enhancements in detail.

> **Note**
>
> To use the advanced mode in Veeam Intelligence, version 13.0.1 or later must be installed on your system. It is not available with the community edition or NFR versions of Veeam.

Malware threat analysis agent

The new **malware threat analysis agent** represents a significant advancement in backup security intelligence. This AI agent uses machine learning to continuously scan workloads at rest and in motion, detecting suspicious changes, validating clean restore points, and modeling the blast radius of detected threats. When anomalies are identified, the agent summarizes the findings, highlights the potential impact per workload, and proposes specific remediation steps.

Operators can review relevant log snippets directly in Veeam Intelligence, identify the nearest clean restore point for affected workloads, and, optionally, trigger signature-based malware scans without leaving the interface. This integration dramatically improves triage quality and accelerates safe recovery decisions by providing all necessary context in a single conversational flow.

Technet24

Deep data analysis agent

The **deep data analysis agent** continuously analyzes backup health and performance, generating insights and automating reports that would traditionally require manual data gathering across multiple interfaces. In Veeam ONE environments, this agent enables users to query any saved report in natural language and work with millions of rows of data without rebuilding reports or constructing complex queries.

The agent loads report results into an internal data frame for instant filtering, compact table generation, and on-demand chart creation. Administrators can ask questions like "*Show protected versus unprotected workloads for the last six months*" and receive visual answers they can export for presentations or compliance documentation. This capability reuses existing report permissions, ensuring security boundaries are maintained while accelerating day-2 analysis tasks.

Data resilience daily summary (morning coffee report)

Version 13 introduces the **data resilience daily summary**, commonly referred to as the **morning coffee report**. This AI-powered email delivers a comprehensive snapshot directly to administrators' inboxes each morning, providing a prioritized overview of job health, detected anomalies, backup status, and recommended action items.

The report aggregates job outcomes across large estates, groups errors by workload for instant triage, and adds context-aware recommendations with links to relevant Veeam knowledge base articles and guides. For organizations managing hundreds or thousands of protected workloads, this daily briefing reduces alert fatigue, shortens morning stand-up preparation time, and gives leaders a clear snapshot they can act on immediately.

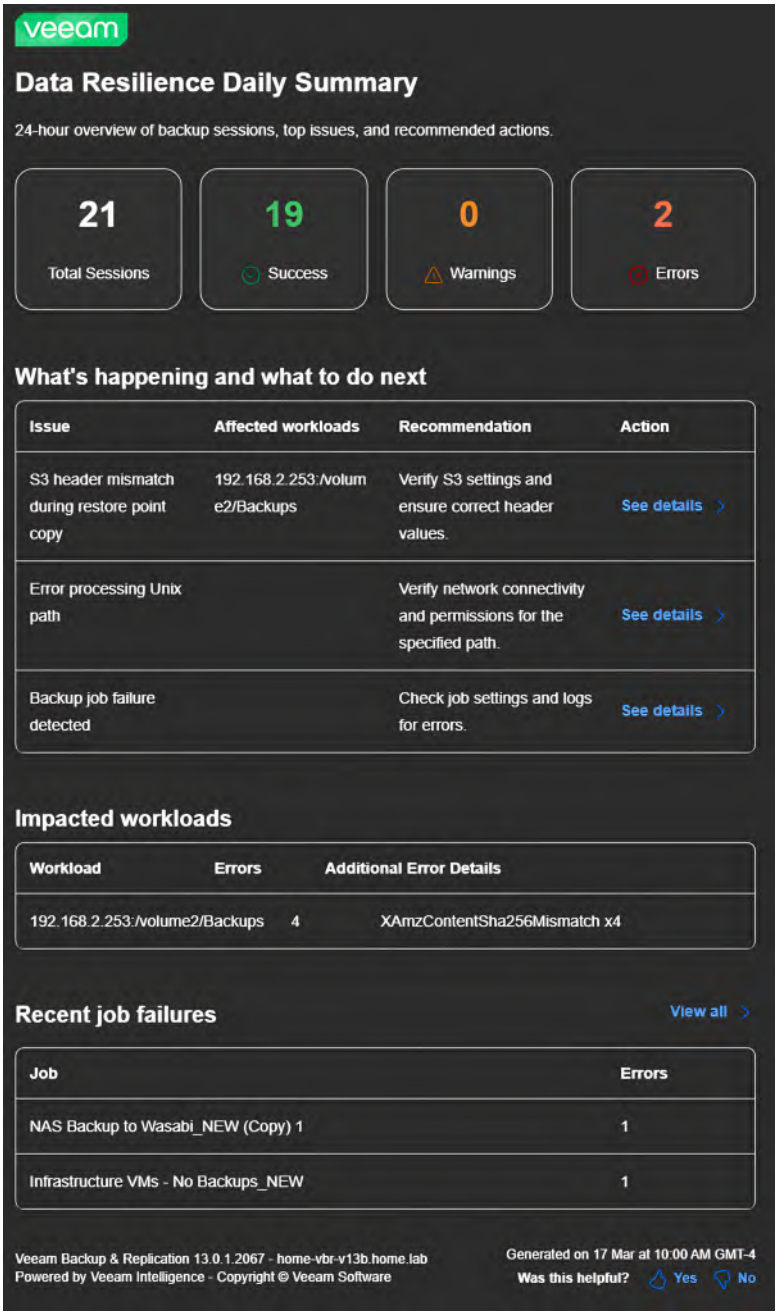


Figure 5.3: Morning coffee report example

Technet24

Enhanced visualization and reporting

Veeam Intelligence in v13 gains significant visualization capabilities that transform how data is presented. The system can now render donut, pie, bar, column, and line charts directly from advanced data sources including Veeam ONE saved reports, alarms, protected workloads, and capacity datasets. This speeds trend reviews, highlights anomalies visually, and improves executive reporting without requiring administrators to leave the product.

The interface also features an expandable modal that provides wide layout options for reviewing long tables, deep multi-step guidance, and side-by-side context. This reduces scrolling and helps operators act faster when reviewing complex answers from AI analysis.

Voice interaction with speech-to-text

Version 13 introduces voice interaction capabilities through integrated **speech-to-text** functionality. Users can ask questions by voice, observe a live waveform while speaking, and receive accurate transcripts that drive the same environment-aware answers as typed queries. This feature improves accessibility and proves particularly valuable during hands-busy workflows such as rack-side troubleshooting or cleanroom recovery operations.

Transparent AI processing

To improve transparency and trust in AI-generated responses, Veeam Intelligence now displays its thought process as a step-by-step task list. This includes key operations such as data retrieval, environment setup, analysis, and response generation. By showing how conclusions are reached through the agentic data pipeline, administrators can better understand and validate the recommendations they receive.

Summary of version 13 enhancements

Table 5.1 summarizes the key enhancements in Veeam Intelligence v13:

Feature	Description
Advanced mode	Unlocks specialized AI agents for malware analysis and deep data analysis
Malware threat analysis agent	ML-powered threat detection with blast radius modeling and clean restore point identification
Deep data analysis agent	Natural language querying of Veeam ONE reports with millions of data rows

Feature	Description
Morning coffee report	AI-generated daily email summarizing job health, anomalies, and prioritized actions
Data visualization	On-demand chart generation, including donut, pie, bar, column, and line charts
Speech-to-text	Voice interaction capability with live waveform feedback for hands-free operation
Transparent AI processing	Step-by-step display of AI processing to improve transparency and trust
Expandable interface	Wide layout modal for reviewing complex tables and multi-step guidance

Table 5.1: Summary of version 13 enhancements for Veeam Intelligence

Benefits of using Veeam Intelligence

Implementing Veeam Intelligence delivers measurable improvements across operational efficiency, security posture, and decision-making speed. We will now discuss the key benefits organizations can expect when leveraging these AI-powered capabilities in detail.

Accelerated troubleshooting and root cause analysis

Traditional backup troubleshooting often requires navigating multiple interfaces, correlating log files, and consulting documentation to identify the source of failures. Veeam Intelligence eliminates this friction by enabling administrators to describe problems in natural language and receive contextualized guidance based on their specific environment. The system identifies early warning signs, uncovers root causes, and provides one-click guided fixes that connect detected issues to trusted solutions.

This capability proves particularly valuable for organizations with distributed teams or varied skill levels among backup administrators. Junior staff can resolve issues that would previously require escalation, while senior engineers can focus on strategic initiatives rather than routine troubleshooting.

Proactive risk management

Veeam Intelligence centralizes risk signals in a unified scoring view based on organizational parameters. Rather than scattering security alerts and backup health indicators across disparate tools, the system correlates findings to provide prioritized response paths. This AI-

Technet24

powered correlation helps security teams focus on the most critical alerts, cutting through noise and reducing time spent on false positives.

The integration with Veeam's threat detection capabilities, including Recon Scanner 3.0 and the malware threat analysis agent, provides early detection of potential threats within restore points. Backup administrators gain visibility into security risks that might otherwise go undetected until recovery time, when the consequences would be most severe.

Simplified compliance and reporting

Compliance requirements demand documented proof of data protection controls, recovery testing, and security measures. Veeam Intelligence simplifies compliance by generating AI-driven reports that demonstrate control effectiveness and meet regulatory requirements. All insights and recommendations are compiled into unified, audit-ready reports that can be exported as needed.

The system's comprehensive audit trail logs every interaction, ensuring organizations can demonstrate appropriate governance over AI-assisted operations. Combined with Veeam's encryption, immutability, and role-based access controls, these safeguards ensure environments remain secure, auditable, and compliant with evolving AI-focused regulatory frameworks.

Improved cross-team communication

One of the most significant benefits of Veeam Intelligence is its ability to bridge communication gaps between technical teams and business stakeholders. Risk scores are tied to business impact, enabling executives, security teams, and backup administrators to communicate using a shared language of risk rather than technical metrics that may not translate across organizational boundaries.

Both technical and non-technical users benefit from clear responses that help them understand their environment. The morning coffee report provides leadership with daily visibility into data protection status without requiring deep technical expertise, while detailed analysis capabilities remain available for engineering teams requiring granular investigation.

Reduced threat dwell time

By surfacing anomalies directly in the data protection layer, organizations can spot attacks earlier, often before they spread into production systems. The malware threat analysis agent's ability to model blast radius and identify clean restore points means recovery planning can begin immediately upon threat detection rather than after lengthy forensic investigation.

This reduction in dwell time directly impacts recovery time objectives and minimizes the potential business impact of security incidents. Organizations can move from detection to validated recovery in a fraction of the time required with traditional approaches.

Enhanced operational efficiency

AI-powered automation continuously monitors system health, prioritizes alerts, and summarizes key insights into daily reports. This removes noise from operations, automates policy tuning, and provides clear visibility into backup performance without requiring complex dashboards or manual status checks.

The deep data analysis agent's ability to generate observability reports instantly means administrators spend less time searching for data and more time making decisions. Natural language querying eliminates the need to understand complex query syntax or navigate nested menu structures, making powerful analysis capabilities accessible to the entire team.

Practical use cases

Organizations across industries are leveraging Veeam Intelligence to solve specific operational challenges. The following examples illustrate common implementations:

- **Financial services**: Banks use Veeam Intelligence to monitor anomalies and guarantee clean recovery of critical banking workloads, supporting regulatory compliance requirements.
- **Healthcare**: Hospitals use AI-powered verification to detect backup tampering, ensuring patient data integrity and compliance with healthcare regulations.
- **Manufacturing**: Logistics companies monitor backup health and proactively surface anomalies, improving decision speed and reducing operational risk.
- **Managed service providers**: MSPs use intelligence capabilities to differentiate their offerings with built-in threat detection and behavioral analytics that previously required separate security tools.

Veeam Intelligence in version 13 represents a significant evolution in how organizations interact with their data protection infrastructure. By combining specialized AI agents for threat analysis and deep data analysis with natural language interfaces and automated reporting, Veeam delivers capabilities that transform backup management from reactive firefighting to proactive resilience management. We will now look at how to access the Veeam Intelligence chat in both the Veeam console and WebUI.

Accessing Veeam Intelligence

Veeam Intelligence is an AI-powered assistant integrated directly into Veeam Backup and Replication v13, designed to help administrators with common issues, questions, and operational guidance. Built on GPT-4o technology and trained on Veeam's comprehensive technical documentation, Veeam Intelligence provides accurate, context-aware answers in

Technet24

natural language, making it an invaluable resource for both day-to-day operations and troubleshooting scenarios.

Prerequisites

Before using Veeam Intelligence, ensure your environment meets the following requirements:

- **License requirements:** A paid license with an active support contract, or an Evaluation license is required. Veeam Intelligence is not available for community (free) edition or NFR licenses.
- **Network connectivity:** The backup server must have an active internet connection with access to the following endpoints:
  - `rest-ai.veeam.com`
  - `cdn.rest-ai.veeam.com`
  - `js.veeam.com`
- **Usage limits:** Each license is limited to 200 questions per 24-hour period. Once this limit is reached, Veeam Intelligence will not process additional queries until the 24-hour period resets.
- **Supported content:** Screenshots and images are not supported in chatbot prompts. All queries must be text-based.

Accessing Veeam Intelligence in web UI

The Veeam Backup & Replication Web UI provides the primary interface for accessing Veeam Intelligence in modern v13 deployments.

Launching Veeam Intelligence

1. Log in to the Veeam Backup and Replication Web UI by navigating to `https://<backup-server>:9443` in your web browser.
2. Once authenticated, locate the **Veeam Intelligence** button in the upper right corner of the interface. The button displays an AI assistant icon.
3. Click the **Veeam Intelligence** button to open the chat window.

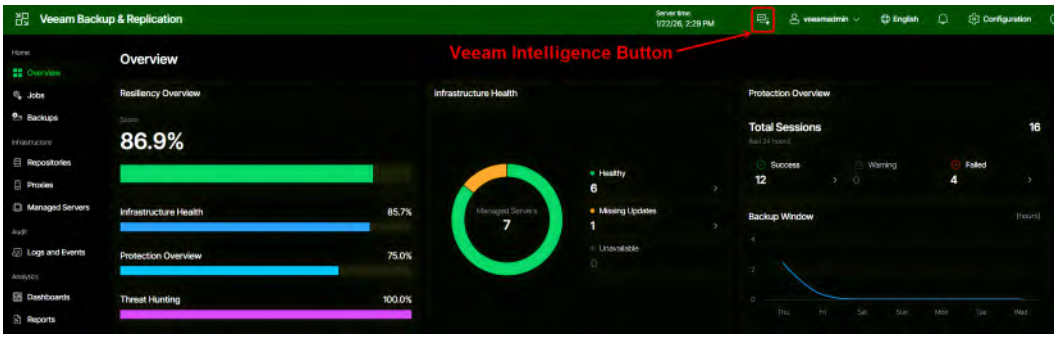


Figure 5.4: Accessing Veeam Intelligence in Web UI

Using the chat interface

The Veeam Intelligence chat window provides an interactive interface for submitting queries:

1. In the chat window, type your question in natural language. You can ask questions in any language, and the assistant will respond accordingly. Here are a few example queries:
    - "*How can I use Veeam Backup & Replication to make sure my VMs are regularly backed up, easily recovered, and securely stored?*"
    - "*What are the best practices for configuring immutable backups?*"
    - "*Help me troubleshoot a failed backup job*"
    - "*How do I add a new vCenter server?*"
2. Click the send button or press **Enter** to submit your question.

Technet24

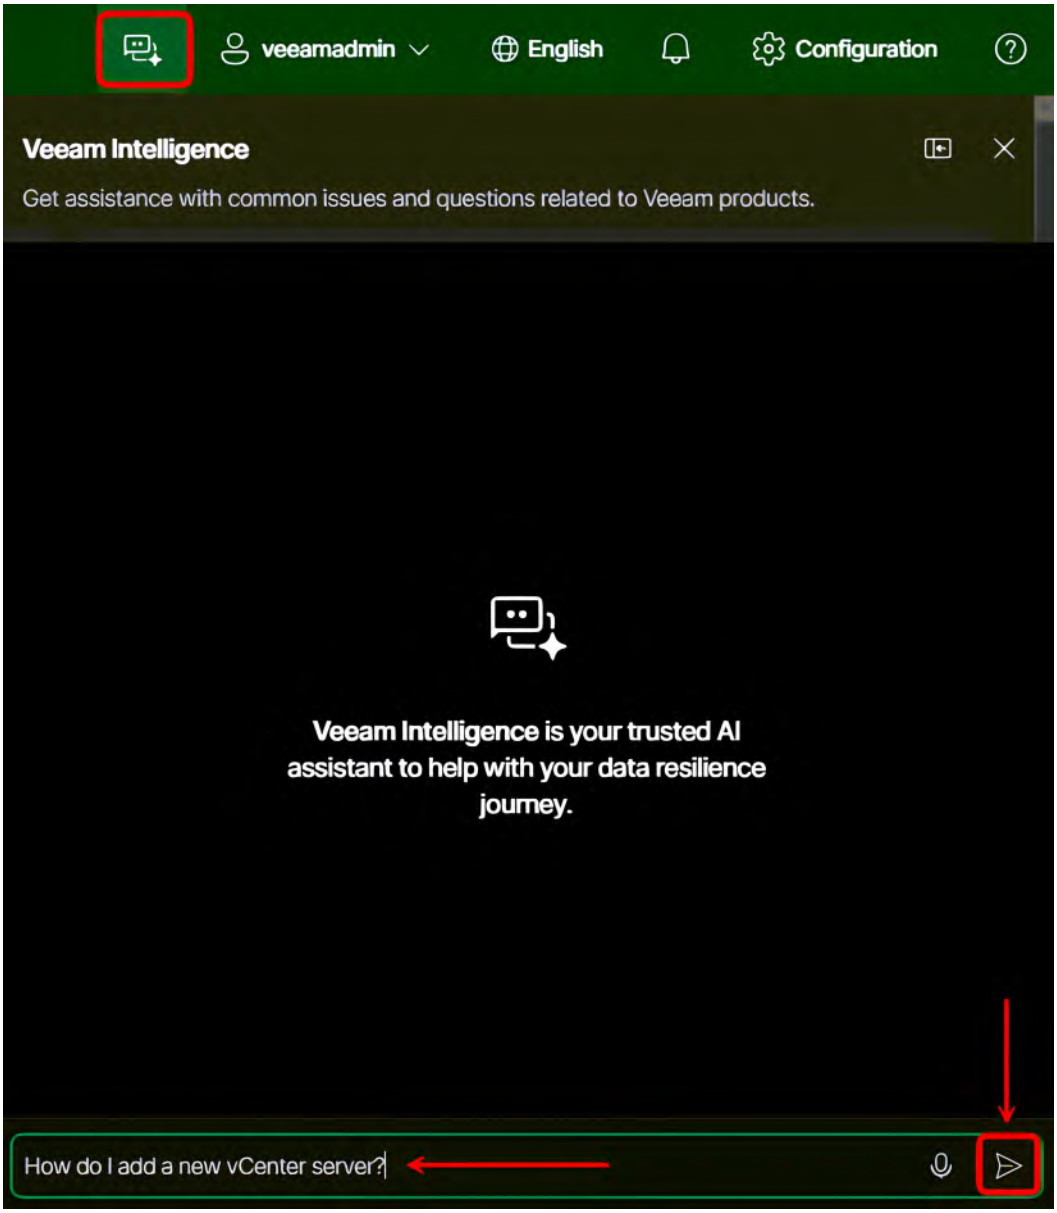


Figure 5.5: Veeam Intelligence chat interface

Responses are generated in real-time, token-by-token, allowing you to begin reading as the AI continues generating the answer. Each response includes relevant links to Veeam documentation and knowledge base articles integrated within the text.

3. If the initial answer is insufficient, add more details to your follow-up question. Veeam Intelligence retains conversation context within the current session, so you do not need to repeat previously shared information.

Configuring Veeam Intelligence settings via web UI

Follow these steps to configure Veeam Intelligence mode or disable the feature entirely:

1. In the Veeam Backup & Replication Web UI, click **configuration** in the top bar (gear icon.)
2. In the configuration menu, click **Veeam Intelligence**.

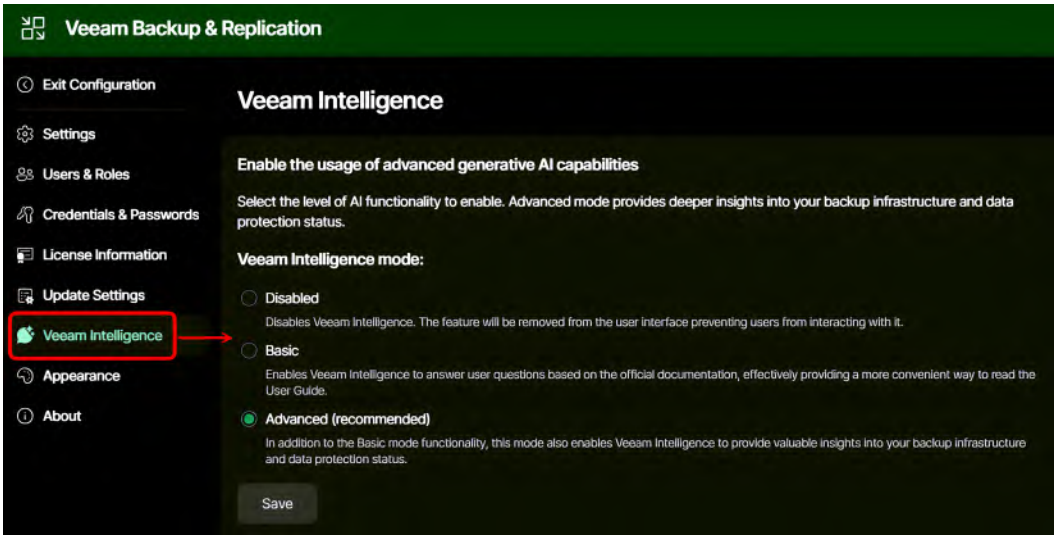


Figure 5.6: Veeam Intelligence settings in configuration section

3. Under **Veeam Intelligence mode:**, select one of the following options:
    - **Disabled**: Disables Veeam Intelligence and hides it from the user interface entirely.
    - **Basic**: Enables Veeam Intelligence in basic mode, providing access to documentation-based assistance.
    - **Advanced**: Enables Veeam Intelligence in advanced mode, providing infrastructure-aware insights and access to specialized AI agents.
    - Click **Save** to save your settings.

Technet24

Accessing Veeam Intelligence in Windows console

The traditional Windows-based Veeam Backup and Replication console also provides access to Veeam Intelligence, ensuring administrators can leverage AI assistance regardless of their preferred management interface.

Launching Veeam Intelligence

1. Launch the Veeam Backup and Replication Console and connect to your backup server.
2. Once connected, locate the **Veeam Intelligence** button in the upper right corner of the console window. Click the button to open the Veeam Intelligence chat window.

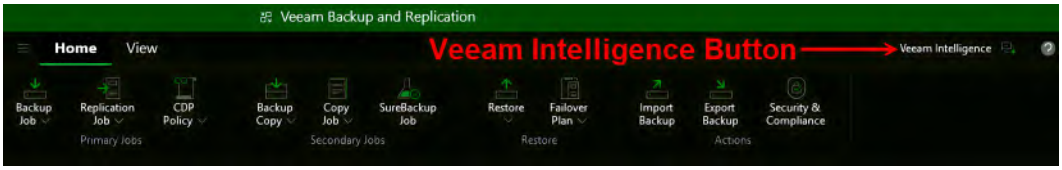


Figure 5.7: Veeam Console accessing Veeam Intelligence

Managing the chat window

The chat window offers several management options:

- **Repositioning:** You can move the chat window and position it anywhere within the UI. The window remains open alongside any wizard you may be working through.
- **Minimizing:** Click the minimize icon to hide the chatbot window while preserving the conversation in the current session.
- **Closing:** Click the close icon to close the Veeam Intelligence window. Note that if you close the session or refresh the page, Veeam Intelligence restarts and loses the context of your previous conversation.
- **Starting new conversation:** Click the **+** button to erase the current conversation and its context, starting fresh.

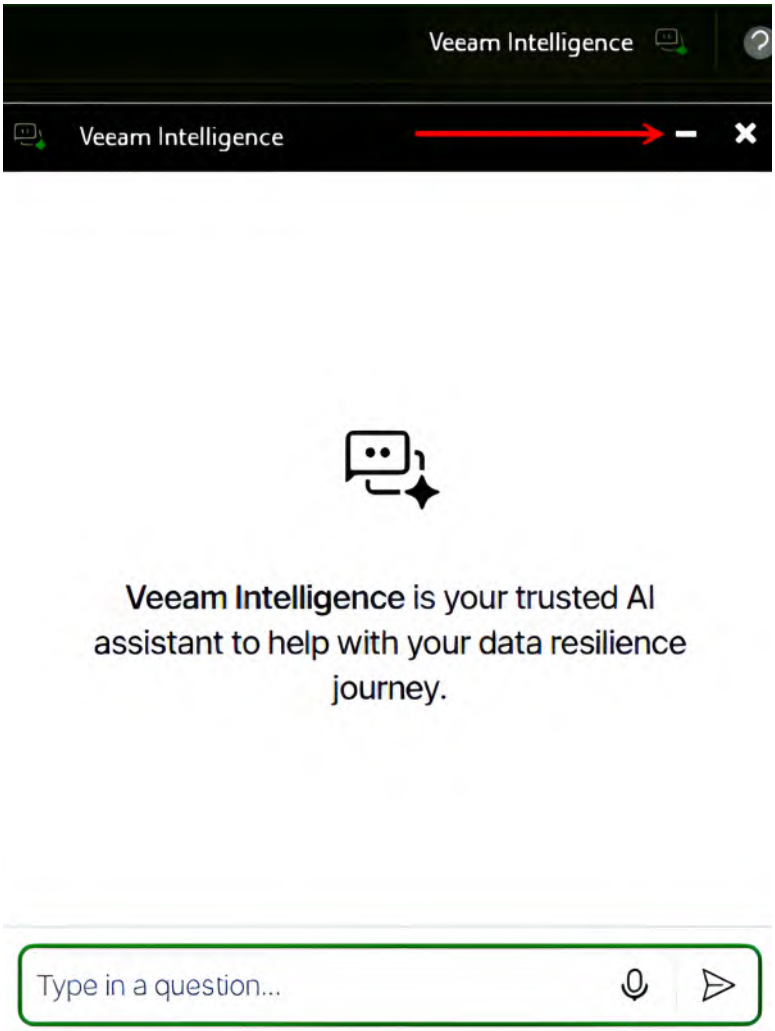


Figure 5.8: Veeam Intelligence window in Veeam console

Using the chat interface

The Windows console chat interface operates identically to the Web UI version. Follow these steps to use the chat interface:

1. Type your question in the chat window using natural language.
2. Click the send icon or press **Enter** to submit your query.
3. Review the response, which includes inline links to relevant documentation.
4. Continue the conversation by adding follow-up questions as needed.

Technet24

Configuring Veeam Intelligence settings via console

To configure Veeam Intelligence settings in the Windows console, follow these steps:

1. From the main menu, navigate to **Options** or access the application settings.
2. Locate the **Veeam Intelligence** settings section and select your preferred mode:
    - **Disabled**: Completely disables the feature
    - **Basic**: Documentation-based assistance only
    - **Advanced**: Full infrastructure-aware capabilities
3. Save and apply the changes.

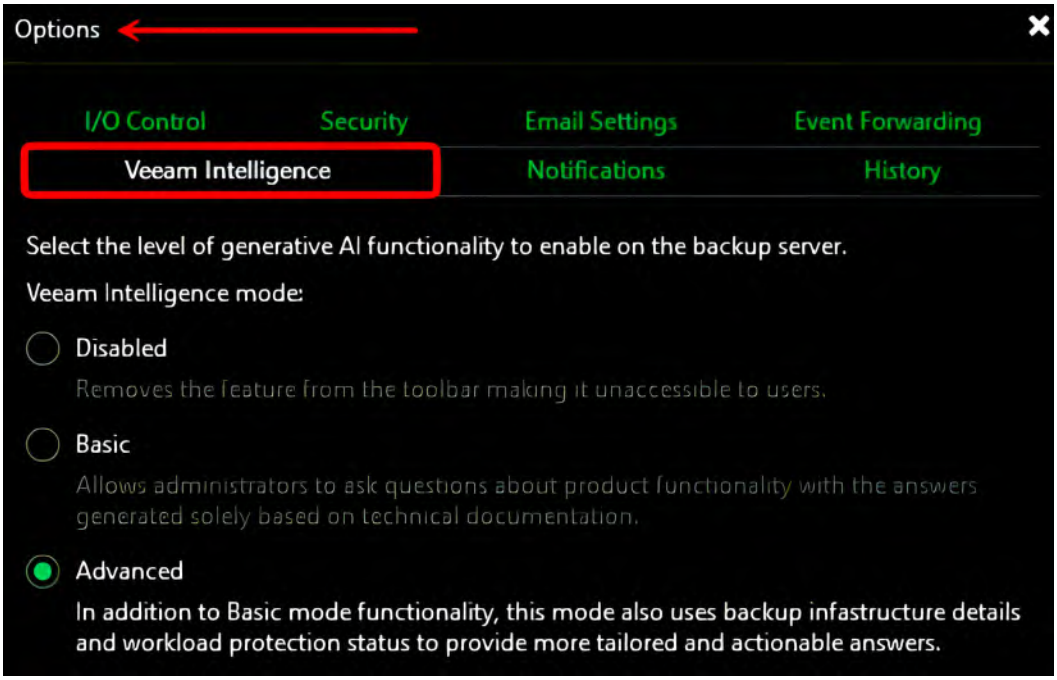


Figure 5.9: Veeam Intelligence settings in console

Best practices

When working with Veeam Intelligence, consider the following recommendations:

- **Be specific**: While Veeam Intelligence can handle both simple and complex queries, providing specific details about your environment, error messages, or desired outcomes will yield more accurate and actionable responses.
- **Leverage context**: Take advantage of conversation context by building on previous questions rather than starting new conversations for related topics.

- **Review documentation links**: Always review the documentation and knowledge base (KB) article links provided in responses for additional context and official guidance.
- **Use advanced mode for operations**: If you need assistance with day-to-day operations, troubleshooting specific jobs, or understanding your infrastructure state, ensure advanced mode is enabled for the most comprehensive assistance.
- **Mind the usage limits**: With 200 questions per 24-hour period, prioritize complex queries that benefit from AI assistance over simple questions that could be answered through documentation search.

Security considerations

When using Veeam Intelligence, keep the following security considerations in mind:

- **Data transmission:** Queries are sent to external Veeam AI services for processing. Do not share confidential or sensitive information with your queries.
- **Response accuracy:** Veeam takes no responsibility for the accuracy of information provided by the chatbot. Always validate critical operational decisions against official documentation and best practices.
- **Disabling for compliance:** If your organization's security policies or compliance requirements prohibit the use of external AI services, you can completely disable Veeam Intelligence through the configuration settings in either the Web UI or Windows console.

Troubleshooting

If Veeam Intelligence is not functioning as expected, verify the following:

- **License status**: Confirm you have a valid paid license with active support or an Evaluation license.
- **Network connectivity**: Ensure the backup server can reach the required Veeam AI service endpoints through any firewalls or proxy servers.
- **Feature status**: Check that Veeam Intelligence is enabled (not disabled) in the configuration settings.
- **Usage limits**: If queries are not being processed, you may have reached the 200-question daily limit. Wait for the 24-hour period to reset.
- **Service availability**: If the Veeam Intelligence service is experiencing high load, requests may be temporarily canceled. Retry your query after a brief wait.

We have now looked at how you access the Veeam Intelligence dialog window in both the Web UI and the Veeam console. Now let's discuss the different modes that you can use with Veeam Intelligence.

Technet24

Veeam Intelligent operating modes

With version 13 of the VSA, Veeam Intelligence has been significantly enhanced through a multi-agent architecture and expanded capabilities, transforming how administrators interact with their backup infrastructure. The assistant operates in one of three configurable modes: disabled, basic, or advanced. Understanding these modes and their capabilities is essential for maximizing the value of your Veeam investment while maintaining appropriate security controls.

Understanding the operating modes

Veeam Intelligence provides three distinct operating modes, each offering different levels of functionality and data access. Organizations can choose the mode that best aligns with their operational requirements and security policies.

Mode	Description	Data access
Disabled	Veeam Intelligence is completely turned off and unavailable to users	None
Basic	Provides documentation-based assistance using agent-based real-time searches across Veeam technical resources	Veeam Help Center, Support KB, R&D Forums, and veeam.com content
Advanced	Includes all Basic mode features plus environment-specific insights using live operational data	All Basic sources plus infrastructure details, monitoring data, and API access

Table 5.2: Veeam Intelligence mode comparison and access

Disabled mode

When Veeam Intelligence is set to Disabled mode, the AI assistant functionality is disabled. The Veeam Intelligence button will not appear in the user interface, and no AI-related features will be available to administrators or users.

Organizations may choose to disable Veeam Intelligence for several reasons:

- Company security policies that prohibit AI-based tools or external data transmission.

- Regulatory compliance requirements in highly restricted environments.
- Air-gapped or isolated network environments without internet connectivity.
- Preference for traditional support channels and documentation access methods.

Basic mode

Basic mode is enabled by default when Veeam Intelligence is activated. In this mode, administrators can submit natural language queries about product functionality, and the assistant responds using information gathered from official Veeam technical documentation and publicly available resources.

Agent-based retrieval

With version 13, Basic mode has been enhanced with agent-based retrieval technology. Rather than relying on a static vector database, Veeam Intelligence now performs real-time searches across multiple Veeam resources to ensure responses are based on the most current and comprehensive information available. This approach enables more accurate and reliable answers by dynamically gathering information from multiple sources.

Basic mode searches the following resources in real time:

- Veeam help center and official product documentation
- Support KB articles
- R&D forums and community resources
- Additional content available on the Veeam official website

Basic mode is ideal for scenarios where administrators need quick answers to product-related questions without requiring environment-specific context. This mode helps you to understand how to configure specific Veeam features, troubleshoot common errors and issues, and learn best practices for backup job configuration. It also helps with finding documentation for specific PowerShell cmdlets or REST API endpoints and getting guidance on upgrade procedures and compatibility requirements.

Advanced mode

Advanced mode builds upon Basic mode by integrating with your backup infrastructure to provide environment-aware insights and recommendations. This mode leverages live operational metrics, backup infrastructure details, and workload protection status to deliver tailored, actionable information specific to your deployment.

Multi-agent architecture

Version 13 introduces a sophisticated multi-agent architecture in advanced mode. A dedicated agent handles conversational queries and technical documentation exploration, while additional on-demand agents are provisioned to address specialized data analysis topics. This

Technet24

architecture enables more comprehensive responses that combine documentation knowledge with real-time infrastructure analysis.

Data access in advanced mode

When operating in advanced mode, Veeam Intelligence can query the following data from your environment through REST API GET queries:

- Backup server configuration and installed version information
- Job health status, outcomes, durations, and success rates
- Proxy and repository capacity utilization
- Protected workloads and restore point information
- License status and instance consumption
- Malware events by severity, source, and time with recovery hints
- Protection groups and discovered agents with install policies
- NAS and file server processing configuration
- Session timeline with failure reasons and performance metrics

Malware deep analysis agent

A significant addition in version 13.0.1 is the malware deep analysis agent, which transforms malware detection events into actionable intelligence. This specialized agent summarizes anomalies, highlights the blast radius per workload, and proposes remediation steps. Operators can review relevant log snippets, identify the nearest clean restore point, and optionally trigger signature-based malware scans directly from within Veeam Intelligence. This capability significantly improves triage quality and speeds safe recovery decisions during security incidents.

> **Note**
>
> **Data visualization**:
>
> Advanced mode supports data visualization capabilities which we have already discussed in the chapter.

Code generation and analysis

Both basic and advanced modes support code generation and analysis capabilities. Veeam Intelligence can generate PowerShell scripts and REST API code for backup server management tasks. The assistant also analyzes submitted REST API or PowerShell scripts, identifies issues, and suggests modifications to resolve problems.

The following are a few example prompts for code generation:

- "*Write a Python script to get information about my failed backup jobs*"
- "*Create a PowerShell script to list all VMs without recent restore points*"
- "*Generate REST API code to retrieve malware detection events after a specific date*"

The enhanced Python code executor in the Veeam Intelligence sandbox supports advanced requests, including sorting results and combining or filtering multi-API outputs, to improve data processing.

Veeam Intelligence: AI-powered assistance for backup administrators

The role of a backup administrator has grown increasingly complex. Modern data protection environments span on-premises infrastructure, multiple cloud platforms, and countless endpoints—each with unique protection requirements, compliance mandates, and recovery objectives. Veeam Intelligence represents a significant step forward in addressing this complexity by bringing artificial intelligence directly into the backup administrator's workflow.

Veeam Intelligence is integrated into the Veeam Data Platform that enables administrators to interact with their backup infrastructure using natural language. Rather than navigating through multiple console screens or memorizing PowerShell cmdlet syntax, administrators can ask questions or request actions in plain English. The assistant interprets these requests, queries the appropriate data sources, and returns actionable information or executes the requested operations.

At its core, Veeam Intelligence serves as an intelligent interface layer that sits between the administrator and the underlying backup infrastructure. It understands context, can correlate information across multiple data sources, and presents findings in a format that facilitates rapid decision-making. This approach fundamentally changes how administrators interact with their Veeam environment, shifting from a tool-centric model to a conversation-centric one.

Simplifying daily operations

The daily routine of a backup administrator typically begins with reviewing the previous night's backup results. Traditionally, this involves opening the Veeam Backup and Replication console, navigating to the Home view, examining session statistics, filtering for failures, and drilling into individual job logs to understand what went wrong. With Veeam Intelligence, this entire workflow condenses into a single query.

Technet24

An administrator can ask "*What failed last night?*" and receive an immediate summary of all failed backup jobs, including the affected workloads, the reasons for failure, and suggested remediation steps. The assistant correlates error codes with known issues and can recommend specific actions based on the failure pattern. If a particular VM failed due to a VSS timeout, Veeam Intelligence might suggest increasing the VSS preparation timeout or checking for known application conflicts on that specific guest.

This natural language interface proves particularly valuable when investigating issues that span multiple areas of the infrastructure. Consider a scenario where backup performance has degraded over the past week. An administrator can ask, "*Why are my backups running slower this week?*" and the assistant will analyze job duration trends, repository throughput metrics, source-side read performance, and network utilization to identify the most likely bottleneck. This type of root cause analysis would traditionally require examining multiple reports and mentally correlating data points—a process that Veeam Intelligence automates and accelerates.

Capacity planning and forecasting

One of the more time-consuming aspects of backup administration involves capacity planning. Administrators must continuously monitor repository utilization, forecast growth trends, and ensure sufficient capacity exists to meet retention requirements. Veeam Intelligence transforms this from a periodic manual analysis into an on-demand conversation.

Asking "*When will my backup repository run out of space?*" prompts the assistant to analyze historical growth patterns, current utilization levels, and projected data change rates to estimate when capacity will become critical. The response includes not just a date but also context on the assumptions used and factors that could accelerate or delay the projection. Administrators can follow up with questions like "*What if I reduce retention from 30 to 21 days?*" to model different scenarios without manually calculating the storage impact.

This capability also extends to planning for new workloads. Before onboarding a new application to the backup infrastructure, an administrator can describe the workload characteristics and ask Veeam Intelligence to estimate the storage impact. The assistant considers the configured backup mode, the expected change rate based on similar workloads, and the retention policy to determine a realistic capacity requirement.

Compliance and reporting

Regulatory compliance and audit preparation represent another area where Veeam Intelligence delivers significant value. When auditors request evidence of backup testing, recovery point compliance, or data protection coverage, administrators traditionally spend hours generating reports and compiling documentation. Veeam Intelligence can answer these questions conversationally and generate supporting evidence on demand.

An administrator preparing for an audit can ask, "*Show me all VMs that don't have a backup in the last 24 hours*," to immediately identify protection gaps. Follow-up questions like "*Which of these are production systems?*" or "*Does replication cover any of these instead?*" help refine the analysis and ensure accurate reporting. When auditors ask about recovery capabilities, the administrator can request, "*Generate a report of all SureBackup verification results from the past quarter*," to demonstrate that recoverability is regularly validated.

Knowledge assistance and best practices

Beyond operational queries, Veeam Intelligence serves as an always-available knowledge resource. Administrators can ask configuration questions, request best-practice guidance, or seek clarification on feature behavior without leaving the console or searching external documentation.

When configuring a new Scale-Out Backup Repository, an administrator might ask "*What's the recommended block size for a performance tier using ReFS?*" and receive guidance tailored to their specific use case. If uncertain about the implications of a particular setting, they can ask, "*What happens if I enable per-VM backup chains?*" to understand the trade-offs before making a configuration change.

This contextual knowledge assistance proves especially valuable for administrators managing environments they didn't originally design. Rather than reverse-engineering configuration decisions made by predecessors, they can ask Veeam Intelligence to explain the current configuration and its implications.

Practical considerations

While Veeam Intelligence significantly streamlines many administrative tasks, it's important to understand its role as an assistant rather than a replacement for administrative expertise. The assistant excels at surfacing information, identifying patterns, and suggesting actions, but the administrator retains responsibility for validating recommendations and making final decisions.

Complex troubleshooting scenarios may still require traditional diagnostic approaches, particularly when issues involve interactions with external systems or edge cases not well-represented in the assistant's training data. An example of this would be a connection from a repository server to the storage used as the repository drive. Administrators should view Veeam Intelligence as a force multiplier that handles routine inquiries and accelerates initial analysis, freeing them to focus on the complex problems that genuinely require human judgment.

The assistant's effectiveness also depends on the quality of information available in the Veeam environment. Accurate job descriptions, meaningful VM tags, and well-organized repository

Technet24

structures all contribute to more relevant and accurate responses. Administrators who invest in maintaining clean metadata will realize greater benefits from Veeam Intelligence.

Veeam Intelligence represents a meaningful evolution in how backup administrators interact with their data protection infrastructure. By providing a natural language interface to complex backup environments, it reduces the cognitive load associated with daily operations, accelerates troubleshooting, simplifies capacity planning, and streamlines compliance reporting. For administrators managing increasingly complex, distributed environments, this AI-powered assistance translates directly into time savings and reduced operational overhead, allowing them to focus on strategic initiatives rather than routine monitoring and investigation.

Summary

This chapter has provided you with an understanding of what Veeam Intelligence is and what is new in Veeam Backup and Replication v13. We discussed how you can access Veeam Intelligence through both WebUI and the Veeam console application. Afterward, we examined how to configure the modes for Veeam Intelligence and what each one allows you to accomplish. Lastly, we explored how Veeam Intelligence can be used by backup administrators to make their management tasks much easier. In the next chapter, we will cover some best practices when using Active Directory with Veeam Backup & Replication v13.

Further reading

- Veeam Intelligence - `https://helpcenter.veeam.com/docs/vbr/userguide/veeam_ai_online_assistant.html?ver=13`
- Veeam Intelligence Modes - `https://helpcenter.veeam.com/docs/vbr/userguide/veeam_ai_online_assistant.html?ver=13#veeam-intelligence-modes`

6

Best Practices for Active Directory Domains with Veeam

Active Directory can be used with Veeam Backup & Replication v13, to leverage segregation and enhanced security. This chapter examines how to use Active Directory with Veeam Backup & Replication. We will review best practices for using Active Directory with your Veeam installation. Lastly, we will examine how to secure your Veeam installation when using it with Active Directory.

In this chapter, we're going to cover the following main topics:

- Active Directory integration with Veeam Backup & Replication
- Best practices for Active Directory with Veeam Backup & Replication v13
- Securing Active Directory for Veeam Backup & Replication v13

Active Directory integration with Veeam Backup & Replication

One of the most significant operational advantages of deploying Veeam Backup & Replication in enterprise environments is its native integration with **Microsoft Active Directory (AD)**. Rather than maintaining a separate user database with its own credential lifecycle, organizations can leverage their existing identity infrastructure to authenticate users and authorize access to backup operations. This integration simplifies administration, strengthens security through centralized policy enforcement, and provides a familiar authentication experience for IT staff.

Technet24

Understanding the authentication model

Veeam Backup & Replication v13 supports two distinct authentication methods: local Windows accounts on the backup server and **Active Directory** domain accounts. While local accounts may suffice for small deployments or isolated test environments, domain integration becomes essential as organizations scale their backup infrastructure or require delegated administration across multiple teams.

When a user connects to the Veeam Backup & Replication console, the server validates credentials against the configured authentication source. For domain accounts, this validation occurs through standard Kerberos or NTLM protocols, depending on your environment configuration. The backup server itself must be either domain-joined or have appropriate trust relationships established to authenticate domain users successfully.

> **Note**
>
> Veeam Backup & Replication does not require schema extensions or modifications to your Active Directory environment. The integration relies entirely on standard AD functionality, specifically security group membership and user authentication, making deployment straightforward and minimally invasive to your directory services infrastructure.

Prerequisites for Active Directory integration

Before configuring Active Directory authentication, several prerequisites must be in place. The Veeam Backup & Replication server should be joined to the domain where your administrative users reside or to a domain with appropriate trust relationships. If you are deploying the Veeam Software Appliance (VSA), the appliance must first be joined to the domain through Web UI or command line interface before AD integration becomes available.

Network connectivity between the backup server and domain controllers must permit standard AD communication. This includes the following:

- DNS resolution for domain controller discovery
- Kerberos traffic on TCP/UDP port 88
- LDAP communication on TCP port 389 or LDAPS communication on port 636
- Global catalog queries on port 3268 if your environment spans multiple domains

Firewall rules should permit these connections from the backup server to all domain controllers that might serve authentication requests.

The service account under which Veeam Backup & Replication runs plays a critical role in AD integration. While the service can run under local system for basic functionality, using a domain service account provides additional flexibility for accessing network resources and simplifies troubleshooting of permission-related issues. This service account does not require domain admin privileges. Membership in the local Administrators group on the backup server, combined with appropriate permissions on protected infrastructure, is sufficient.

Configuring role-based access control with AD groups

Veeam Backup & Replication implements a sophisticated role-based access control (RBAC) model that maps automatically to Active Directory security groups. Rather than assigning permissions to individual users, best practice dictates creating AD security groups that correspond to your desired permission sets, then assigning Veeam roles to those groups. An example of a security group would be something like *SG_Veeam_Administrators*. This approach provides several operational benefits: new team members automatically inherit appropriate permissions when added to the relevant groups, departing employees lose access immediately upon group removal, and your backup infrastructure permissions remain synchronized with your broader identity governance processes.

The product has several predefined roles that address common administrative scenarios by default. The **Veeam backup administrator** role grants full access to all features and configuration options, appropriate for senior backup administrators who require unrestricted control. The **Veeam backup operator** role permits execution of existing jobs and basic monitoring but restricts configuration changes, making it suitable for operations staff who manage day-to-day backup activities without modifying the infrastructure. The **Veeam restore operator** role focuses specifically on recovery operations, allowing users to restore data without access to backup configuration or job management.

To configure role assignments using Active Directory groups, follow these steps:

1. Open the Veeam Backup & Replication console and navigate to the **Users and Roles** section under the main menu.
2. Click the **Add** button to add your security principals and browse your Active Directory structure.
   When selecting principals, choose the appropriate AD security groups rather than individual user accounts.
3. After adding a group, assign the desired role from the **Role** dropdown menu. The configuration takes effect immediately. Users who are members of the assigned groups can connect to the console using their domain credentials.

For organizations with complex administrative requirements, the product supports granular scope-based permissions. In addition to assigning a role, you can restrict that role to specific

Technet24

backup infrastructure components. For example, you might grant a team backup operator permissions but limit their scope to only those backup jobs and repositories relevant to their business unit. This capability enables secure multi-tenant administration within a single Veeam deployment, with Active Directory groups defining who has access and scope assignments defining what they can access.

Service account configuration and delegation

Beyond console access, AD accounts serve critical functions throughout the Veeam Backup & Replication infrastructure. The credentials stored in the **Credentials Manager** enable Veeam to authenticate against protected systems, access network shares, and communicate with virtualization platforms. Proper configuration of these service accounts directly impacts both security and operational reliability.

When adding credentials to the **Credentials Manager**, you must specify domain accounts in the standard `DOMAIN\username or username@domain.com format`. The stored credentials are encrypted and cached locally on the backup server. Veeam will use these credentials when connecting to systems that require authentication, such as VMware vCenter servers, Hyper-V hosts, or file share repositories.

For environments with strict security requirements, consider implementing dedicated service accounts for different infrastructure tiers. A separate account for virtualization platform access, another for repository access, and a third for application-aware processing allow precise control over the permissions granted to each function. This separation also simplifies audit trails and reduces the blast radius if any single credential becomes compromised.

Active Directory group managed service accounts (gMSA) provide an additional security enhancement for Veeam deployments. These gMSA accounts feature automatic password rotation managed by the domain, eliminating the operational burden of manual credential updates while strengthening security. Veeam Backup & Replication v13 supports gMSA for the main service account and for credentials used to access Windows-based infrastructure. Implementing gMSA requires a domain functional level of Windows Server 2012 or higher and appropriate key distribution service configuration in your domain.

Security considerations and best practices

Integrating your backup infrastructure with Active Directory extends your identity security perimeter to include backup operations. This extension requires careful attention to several security considerations, like the following, to prevent backup systems from becoming a weak link in your security architecture:

- Protect accounts with console access using multi-factor authentication where possible. While Veeam Backup & Replication does not natively implement MFA for console

logon, you can enforce MFA at the network level through solutions like remote desktop gateway with Azure MFA or by deploying the console only on privileged access workstations subject to conditional access policies.

- Audit logging captures all authentication attempts and administrative actions performed through the console. These logs should be forwarded to your SIEM or central logging platform for monitoring and retention. Pay particular attention to failed authentication attempts, role assignment changes, and credential manager modifications as these events may indicate reconnaissance or privilege escalation attempts.
- Implement the principle of least privilege rigorously. Backup administrators often require significant permissions across the infrastructure to perform their duties, making their accounts attractive targets for attackers. Use separate accounts for daily administrative work versus elevated backup administration tasks and avoid granting backup-related AD groups membership in highly privileged roles like domain admins.
- Consider the implications of backup data access when designing your role assignments. Users with restore permissions can potentially access data from any protected system, regardless of their normal access rights in that production environment. Restrict restore scope to appropriate systems to help maintain data confidentiality even as you delegate recovery operations to distributed teams.

This section has covered some best practices for security considerations. Let's now look at some common integration issues and how to avoid them.

Troubleshooting common integration issues

Authentication failures in AD-integrated environments typically stem from a few common causes:

- DNS misconfiguration prevents the backup server from locating domain controllers, manifesting as timeout errors during login attempts. Verify that the backup server uses domain DNS servers and can resolve the domain name to appropriate domain controller addresses.
- Clock synchronization issues cause Kerberos authentication failures when the time difference between the backup server and domain controllers exceeds five minutes. Ensure that all systems synchronize time from a common authoritative source, typically your domain controllers, which in turn should synchronize with external time sources.
- Group membership changes may not reflect immediately in user sessions. Windows caches security group membership in user tokens, so recently modified group

Technet24

assignments may not take effect until the user logs out and back in or, in some cases, until the token is refreshed through the normal Kerberos ticket lifecycle.

- Verify both the role assignment in Veeam and the user's effective group membership if users can authenticate but lack expected permissions. Nested groups function correctly with Veeam role assignments, but complex nesting can make troubleshooting difficult. The `whoami /groups` command run from the user's context displays effective group membership and can help identify whether the expected groups appear in the user's token.
- Reference the Windows event log on the backup server for service account issues, as it contains detailed authentication and authorization events (like login failures). The security log records logon events, while the system log may contain service-specific errors. Additionally, the Veeam Backup & Replication logs in `C:\ProgramData\Veeam\Backup` provide application-level details about authentication operations and credential usage.

We have now covered some common integration troubleshooting steps that can be performed when the noted issues arise.

Integration with Veeam software appliance

VSA introduces specific considerations for Active Directory integration. The appliance runs on a hardened Linux operating system rather than Windows, which affects how domain integration functions at the operating system level. However, the Veeam Backup & Replication services running within the appliance support the same AD authentication capabilities as traditional Windows deployments.

To enable Active Directory authentication on a software appliance deployment, you must first configure domain membership through the appliance management interface. The Web UI provides a dedicated section for identity configuration where you specify the domain name, provide credentials with permission to join computers to the domain, and optionally designate the organizational unit for the computer object. The appliance uses system security services daemon (SSSD) and Samba integration to communicate with Active Directory, supporting both Kerberos and NTLM authentication methods.

After domain join, configure console user access through the normal Veeam Backup and Replication console by adding AD groups and assigning roles as described earlier. The console authentication process remains identical from the user perspective; the underlying mechanics of validating credentials against Active Directory are handled transparently by the appliance infrastructure.

Time synchronization on the software appliance deserves particular attention. Linux systems do not automatically synchronize with domain controllers the way domain-joined Windows systems do. Configure the appliance to use NTP with your domain controllers as time sources or point to the same external time sources your domain infrastructure uses. The appliance Web UI includes time configuration options, or you can configure NTP through the console using standard Linux tools.

Active Directory integration transforms Veeam Backup & Replication from a standalone application with its own user management into a component of your broader enterprise identity infrastructure. By leveraging AD groups for role-based access control, domain accounts for service authentication, and your existing identity governance processes for access management, you achieve both operational simplicity and an improved security posture.

The key principles for successful AD integration include planning your security group structure to match your administrative model, implementing least privilege through appropriate role and scope assignments, protecting privileged accounts with additional security controls, and maintaining comprehensive audit logging of backup administrative activities. With these foundations in place, your Veeam infrastructure inherits the security investments you have already made in your Active Directory environment while providing flexibility to delegate backup operations safely across your organization.

Now let's take a look at some of the best practices you should follow when using Active Directory with Veeam.

Best practices for Active Directory with Veeam Backup & Replication v13

Active Directory is the identity backbone of nearly every enterprise environment. As organizations deploy Veeam Backup & Replication v13, the relationship between the backup infrastructure and Active Directory requires careful planning across multiple dimensions: protecting domain controllers as critical workloads, leveraging AD for authentication and access control within the Veeam infrastructure, and deciding whether the Veeam components themselves should be domain-joined. Version 13 introduces significant enhancements to identity management, including SAML 2.0 Single Sign-On, enhanced RBAC, and the Linux-based VSA architecture, all of which change the calculus of how Active Directory integrates with your data protection platform.

This section provides best practice guidance for each of these areas, helping you design a resilient, secure, and operationally efficient Veeam environment that properly accounts for Active Directory dependencies and leverages its capabilities where appropriate.

Technet24

Veeam infrastructure and Active Directory: Domain placement strategy

Microsoft Active Directory is the heart of the IT infrastructure for nearly every organization. When setting up the Veeam Data Platform infrastructure, keep in mind the foundational principle that a data protection system should not rely on the environment it is designed to protect in any way. When your production environment goes down along with its domain controllers, any dependency on those controllers for backup console authentication, DNS name resolution, or credential validation will directly impact your ability to perform actual restores at the moment you need them most.

Furthermore, production domains are the place where you have the highest number of users, and therefore the highest number of potential targets for social engineering or phishing attacks. Production domains should be considered inherently more vulnerable than isolated management environments. Once intruders gain access to the production domain, they can find a way to escalate to high privileges. An intruder who gains elevated access to a Veeam Backup & Replication server can extract detailed information about the entire infrastructure it protects. This information allows them to hit harder, potentially driving an attack that has more impact, is more publicly visible, and takes longer to recover from.

This risk is not theoretical. Modern ransomware operators routinely target backup infrastructure as part of their attack chain, and Active Directory compromise is frequently the vector used to reach it.

Understanding Active Directory forests and domains

Before evaluating deployment options, it is helpful to understand the Active Directory structures that underpin each model. Microsoft Active Directory consists of forests and domains.

A **forest** is a collection of one or more Active Directory domains that share a common logical structure, directory schema (class and attribute definitions), directory configuration (site and replication information), and global catalog (forest-wide search capabilities). Domains within the same forest are automatically linked with two-way, transitive trust relationships. This means that any domain in a forest inherently trusts every other domain in that same forest.

A **domain** is a partition within an Active Directory forest. Partitioning data enables organizations to replicate data only to where it is needed, allowing the directory to scale globally over networks with limited available bandwidth. Each domain maintains its own set of security policies, user accounts, and group memberships.

Understanding this distinction is critical because it explains why placing Veeam components in the same forest as your production environment, even in a separate domain, does not

provide true isolation. The transitive trust relationships within a forest mean that a compromise of one domain can be leveraged to attack another domain in the same forest. True isolation requires a separate forest or a workgroup deployment.

Deployment options

When securing administrative accounts and the Veeam Data Platform infrastructure, you have several options ranging from most secure to least secure:

1. Add the Veeam components to a management domain that resides in a separate Active Directory forest and protect the administrative accounts with multi-factor authentication.
2. Add the Veeam components to a separate workgroup and place the components on an isolated network segment where applicable.
3. Add the Veeam components to the production domain but ensure that accounts with administrative privileges are protected with multi-factor authentication.

Note

Best Practice:

For the most secure deployment, add the Veeam components to a management workgroup or a management domain that resides in a separate Active Directory forest, and protect the administrative accounts with MFA. This way, the Veeam Data Platform infrastructure does not rely on the environment it is meant to protect.

Figure 6.1 demonstrates the different options for workgroup- and domain-based configurations for Veeam Backup & Replication v13.

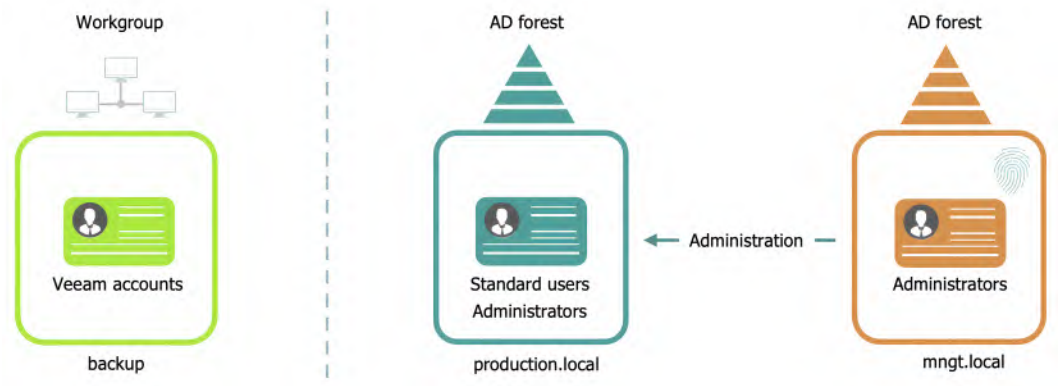


Figure 6.1: Workgroup and domain options

Technet24

Windows workgroup deployment

When using a workgroup, every system must be configured independently with its own local security policy, local user accounts, and permissions. Everything must be carefully documented for management and compliance reasons. If you have multiple Veeam servers and users, this approach can become extremely cumbersome in larger environments. Kerberos authentication is not available with workgroup servers; NTLM is used instead.

It is also important to recognize that a workgroup deployment is harder to defend against insider threats such as a disgruntled employee. Because you are using local accounts on each workgroup server, you cannot simply deactivate a single AD account to lock out that employee from the critical backup infrastructure. Each server would need to have its local accounts individually managed. Additionally, it is harder to prove for compliance reasons that the systems are safe and being utilized as they should be, since there is no centralized policy enforcement or audit trail. A workgroup setup is generally a good solution for small environments.

Workgroup advantages

Here are the advantages of using workgroup setups:

- Fast and easy setup with minimal infrastructure requirements.
- Separates Veeam accounts from domain privileged accounts, which helps protect against keyloggers and credential theft following a breach of the production domain.
- Does not rely on the environment it is meant to protect. Backup operations continue even if all production domain controllers are offline.
- No additional infrastructure servers such as domain controllers, NTP servers, or DNS servers are required.

Workgroup disadvantages

The following are some disadvantages of setting up workgroups:

- Large management overhead in larger environments, as each server requires independent account and policy management.
- No Kerberos authentication when logging into a standalone workgroup server, as only NTLM is available, which is considered less secure.
- Hard to be compliant compared to using a domain, perform compliance checks, and prove compliance to auditors as there is no centralized policy management or audit trail.

- Impossible to use gMSA for guest backup interaction, require traditional password-based service accounts instead.
- Hard to defend against insider threats due to user management, as there is no centralized account that can be deactivated to immediately revoke an individual's access across all Veeam infrastructure servers.

Note

With the introduction of the VSA in v13, the Linux-based VSA naturally operates in a workgroup-like configuration. The VSA does not join an Active Directory domain, which provides inherent isolation from production domain compromise while still allowing AD-based authentication through the Web UI and SAML SSO integration. This makes the VSA an attractive option that combines the security benefits of workgroup isolation with modern identity management capabilities.

Management domain deployment

While a management domain approach does add a forest to the Active Directory environment, the cost and complexity are limited by the fixed design, small hardware and software footprint, and small number of users. Enabling central management of policies, user rights, and permissions makes ongoing administration straightforward. This model also enables one-click deactivation of a single AD account when you face an insider threat, which is a significant advantage over a workgroup deployment.

Setting up a separate forest with a management domain is a great solution for larger environments. You can also integrate MFA on the management domain to protect administrative accounts even further, blocking man-in-the-middle attacks and keyloggers.

Management domain advantages

The following are some of the advantages of the management domain:

- Easy management with centralized policy enforcement and user administration.
- One-click deactivation of a single AD account to immediately revoke access in the event of an insider threat or personnel change.
- Does not rely on the environment it is meant to protect. Operates independently of the production forest.
- Enables secure Kerberos authentication between different Veeam components, providing mutual verification and ticket-based authentication.

Technet24

- Enable control of the domain and fulfillment of compliance requirements more easily, with centralized audit logging.
- Can integrate MFA for an extra layer of security, blocking man-in-the-middle attacks and keyloggers.

Management domain disadvantages

Disadvantages of management domain include the following:

- Requires additional infrastructure components including at least two domain controllers, DNS, and NTP for the management forest.
- Requires more specialized knowledge to set up correctly, including forest trust configuration and proper security hardening.
- While Kerberos is considered more secure than NTLM, it is still vulnerable to attacks such as Kerberoasting, where an attacker extracts service account ticket hashes and cracks them offline. Service accounts in the management domain should use long, complex passwords or gMSA to mitigate this risk.

Deployment model comparison

The following table provides a side-by-side comparison of the three deployment models across key operational and security criteria.

Criteria	Management forest	Workgroup	Production domain
Isolation from production	Full isolation through a separate forest	Full isolation through no domain membership	No isolation, shared domain
Authentication protocol	Kerberos	NTLM only	Kerberos
Centralized account management	Yes (AD-based)	No (per-server local accounts)	Yes (AD-based)
gMSA support	Yes	No	Yes
Group policy enforcement	Yes (management forest GPOs)	No (local security policy only)	Yes (production domain GPOs)
MFA integration	Native AD MFA	Via SAML SSO only	Native AD MFA

Criteria	Management forest	Workgroup	Production domain
Insider threat response	One-click account deactivation	Must disable per-server local accounts individually.	One-click account deactivation
Compliance posture	Strong (centralized audit and policy)	Challenging (per-server documentation)	Strong (but shared with production)
Infrastructure overhead	Moderate (additional DCs, DNS)	None	None
Best suited for	Large enterprises and regulated industries	Small to mid-size environments and VSA deployments	Environments that prioritize operational simplicity

Table 6.1: Feature comparison based on model—workgroup vs domain

Figure 6.2 shows some main pros and cons of the different scenarios with both workgroup and domain.

	Worst practice	Quick win	Best Practice
What	Single Domain	Workgroup	One-way Trusted Management Domain
Pros	Deployment Management Usage	Isolation Deployment	Isolation Management Usage
Cons	First target Pivot for lateral moves	Management Usage	Deployment

Figure 6.2: Pros vs. cons of workgroup and domain

That covers the comparisons and pros/cons of workgroup and domain scenarios. Let us now look at the forest trust configuration requirements.

Forest trust configuration

When deploying a separate management forest, forest trusts are used to manage the segmented Active Directory domain services infrastructure and support access to resources across the forests. Forest trusts allow administrators to connect two AD DS forests with a

Technet24

single trust relationship, providing a seamless authentication and authorization experience across the forests while maintaining administrative autonomy.

The recommended configuration is a **one-way forest trust** where the production domain trusts the management domain. With this arrangement, members of the management forest (Veeam administrators) can utilize resources located in the production forest for backup operations. However, the trust operates in only one direction, so production domain accounts cannot authenticate against or access management forest resources.

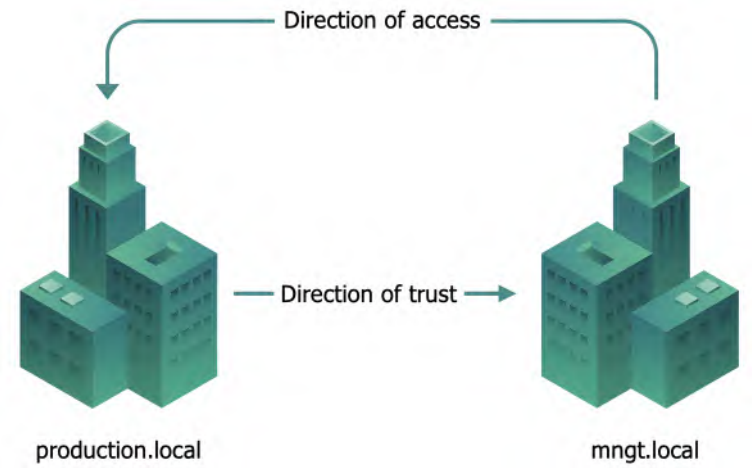


Figure 6.3: Forest trust configuration for security

In practice, this means you will configure an incoming one-way forest trust on the management domain. If the production forest is compromised, the attacker cannot leverage that access to authenticate against the management forest or its member servers, preserving the integrity of your backup infrastructure.

Note

Regardless of which deployment model you choose, always ensure that DNS name resolution is available independently of the production domain. If your Veeam infrastructure relies on production domain controllers for DNS, a domain outage will prevent name resolution and may block restore operations. Consider deploying dedicated DNS servers or configuring conditional forwarders within the management forest.

Authentication and identity management in v13

Veeam Backup & Replication v13 introduces significant enhancements to how users authenticate and access the backup infrastructure. These improvements align with Zero Trust security principles and provide organizations with flexible, enterprise-grade identity management options.

SAML 2.0 single sign-on

Version 13 introduces SAML 2.0-based single sign-on, a major advancement that allows organizations to centralize authentication through external Identity Providers (IdPs). SAML SSO is supported for both the new Web UI and the traditional Windows-based backup console, providing consistent authentication experience regardless of the management interface used.

SAML integration enables organizations to leverage their existing identity infrastructure, including Microsoft Entra ID (formerly Azure Active Directory), Okta, PingFederate, and other SAML 2.0-compliant providers. This delivers several key benefits:

- **MFA passthrough:** MFA policies configured at the IdP level are automatically enforced for Veeam access, eliminating the need to configure MFA separately within Veeam.
- **Centralized authentication:** User access is managed through a single identity provider, reducing password sprawl and simplifying user lifecycle management.
- **Conditional access policies:** Organizations using Entra ID or similar providers can enforce location-based, device-based, and risk-based access policies for Veeam infrastructure access.
- **Streamlined compliance:** Centralized authentication logging and access auditing satisfy regulatory requirements for identity governance.

Tip

When configuring SAML SSO with Microsoft Entra ID, ensure that the Veeam application registration includes the appropriate redirect URIs for both the Web UI and the backup console. Test the authentication flow with a non-administrative account before deploying to production users to verify that group claims and role mappings are correctly configured.

Enhanced role-based access control

Version 13 delivers a significantly enhanced RBAC framework that moves beyond the predefined roles of earlier versions. The new custom role wizard provides a granular interface

Technet24

for defining roles with specific permissions scoped to specific infrastructure objects, job types, or operations.

When integrating RBAC with Active Directory, the following best practices apply:

- Use Active Directory security groups for role assignments rather than individual user accounts.
- Create dedicated AD security groups that map to each Veeam role, such as `Veeam-BackupAdmins`, `Veeam-BackupOperators`, `Veeam-RestoreOperators`, and `Veeam-ReadOnly`.
- Apply the principle of least privilege. Use scoped roles to restrict operators to only the jobs, repositories, or infrastructure objects they need to manage. The v13 custom role wizard makes this significantly easier than in previous versions.
- Remember that MFA in Veeam Backup & Replication requires AD-based accounts. MFA cannot be enabled for local accounts. Additionally, disabling MFA for any account will disable it for all accounts, so exercise caution when modifying MFA settings.

Tip

Important:

When using AD accounts for RBAC, always specify credentials in `DOMAIN\Username` format. For local user accounts, use `Username` or `HOST\Username format`. Inconsistent credential formatting is a common source of authentication failures.

Kerberos vs. NTLM authentication

For domain-joined Veeam components, Kerberos authentication provides stronger security than NTLM by using ticket-based authentication with mutual verification between client and server. Veeam Backup & Replication v13 supports Kerberos authentication for communication between Veeam components when the infrastructure is domain-joined, whether in a production domain or a management forest.

However, Kerberos is not invulnerable. Kerberoasting attacks allow an attacker with a valid domain account to request service tickets for service accounts and then crack those tickets offline to recover passwords. To mitigate this risk in a management domain, use gMSA wherever possible, as their passwords are automatically rotated and are cryptographically complex. Where gMSA cannot be used, enforce long and complex passwords of at least 25 characters for service accounts.

In workgroup deployments or when using VSA, NTLM authentication is used instead. While NTLM is considered less secure than Kerberos, the isolation benefits of a workgroup or VSA deployment typically outweigh the protocol-level security difference, particularly when combined with network segmentation and the hardened operating system of the VSA.

Leveraging Active Directory for agent-based protection groups

Veeam Backup & Replication v13 uses Active Directory integration to dynamically discover and protect physical and virtual machines through protection groups. Rather than manually adding individual computers to backup policies, you can define protection groups based on Active Directory objects such as domains, organizational units, containers, computer objects, failover clusters, or security groups.

Security recommendations summary

The following table consolidates the key security recommendations for Active Directory integration with Veeam Backup and Replication v13.

Area	Recommendation
Domain placement	Deploy Veeam infrastructure in a separate management forest or workgroup. The backup system must not rely on the environment it is meant to protect. Protect administrative accounts with MFA.
Authentication	Implement SAML 2.0 SSO with MFA passthrough for all administrative access. Use Kerberos where available but mitigate Kerberoasting risk with gMSA and complex passwords.
Access control	Leverage the enhanced v13 RBAC framework with AD security groups. Apply least-privilege scoping to restrict operator access to only required resources.
Insider threats	Prefer management domain over workgroup in larger environments to enable one-click account deactivation. In workgroup deployments, document all local accounts and maintain a rapid response plan.
Credential management	Use the centralized Credential Manager. Consider gMSA for guest processing where domain-joined deployments are used. Rotate credentials regularly.

Technet24

Area	Recommendation
DC backup	Back up at least two writable DCs per domain with AAIP enabled. Prioritize FSMO role holders. Never replicate domain controllers with Veeam.
DNS independence	Ensure DNS resolution does not depend solely on production domain controllers. Deploy dedicated DNS or configure conditional forwarders.
Network security	Use LDAPS (port 636) for all directory queries from the Veeam server. Ensure Root CA certificates are properly published.
Backup retention	Align backup retention with the AD tombstone lifetime (default 180 days). Ensure at least one valid backup exists within this window at all times.
Compliance	Choose a deployment model that supports your compliance requirements. Management domains provide centralized audit logging and Group Policy enforcement. Document all accounts and policies in workgroup environments.

Table 6.2: Security recommendations summary

Active Directory integration with Veeam Backup & Replication v13 spans multiple dimensions, from the fundamental question of where your backup infrastructure sits relative to the domain, through the new authentication and access control capabilities, to the specific requirements for protecting domain controller workloads. The overarching principle remains that your data protection system should never rely on the environment it is designed to protect.

The v13 release provides organizations with greater flexibility in how they address this principle. Now let's look at the security aspects and how securing Active Directory should be top of mind.

Securing Active Directory for Veeam Backup & Replication v13

In this section, we will examine the three primary deployment models for positioning Veeam infrastructure relative to Active Directory. You will also get detailed guidance on securing

service accounts, and cover configuration of MFA to harden administrative access in Veeam Backup & Replication v13.

Deployment models: Workgroup, management domain, or production domain

When planning where to place your Veeam Backup & Replication infrastructure relative to Active Directory, three models are available. Each represents a different balance of security isolation, operational complexity, and management overhead. Understanding the trade-offs is essential for selecting the approach that matches your organization's risk tolerance and operational maturity.

The *Veeam Security Best Practices Guide* ranks these options from most secure to least secure as follows:

Rank	Deployment model	Description
1	**Separate management forest**	Veeam components are joined to a dedicated management domain in its own Active Directory forest, with administrative accounts protected by multi-factor authentication.
2	**Workgroup deployment**	Veeam components are configured as standalone workgroup members on an isolated network segment with no Active Directory dependency.
3	Production domain (hardened)	Veeam components remain in the production domain, and administrative accounts are protected by MFA and strict privilege controls.

Table 6.3: Deployment models and rankings

Option 1: Separate management forest (recommended)

This is the most secure deployment model for enterprise environments. A separate Active Directory forest is deployed specifically for management infrastructure, including Veeam Backup & Replication servers, proxies, and any Windows-based repository servers. Because a forest represents the ultimate trust boundary in Active Directory, this design ensures complete isolation from the production environment's authentication and authorization systems.

In this model, a one-way forest trust can be configured so that the production forest trusts the management forest, allowing management accounts to access production resources when necessary for backup and restore operations. Critically, the management forest does not trust

Technet24

the production forest, which means that compromised production credentials cannot be used to authenticate against backup infrastructure components.

The management forest requires its own domain controllers, DNS infrastructure, and optionally its own NTP hierarchy. While this adds infrastructure components, the footprint is deliberately small, typically two domain controllers for redundancy and a handful of management accounts. The operational benefits are significant: centralized group policy enforcement, Kerberos authentication between Veeam components, one-click account deactivation for insider threat response, streamlined compliance auditing, and the ability to integrate multi-factor authentication at the domain level.

Advantages	**Disadvantages**
Complete security isolation from the production environment	Requires additional infrastructure (domain controllers, DNS)
Kerberos authentication between Veeam components	Increased initial setup complexity and AD administration knowledge
Centralized Group Policy management and compliance enforcement	Forest trust configuration requires careful planning.
One-click deactivation of compromised or revoked accounts	Additional licensing considerations for Windows Server domain controllers
Supports domain-level MFA integration for administrative access	Kerberos, while more secure, remains vulnerable to Kerberoasting attacks
Enables gMSA usage for guest processing operations	
Backup infrastructure remains operational if the production AD is compromised	

Table 6.4: Option 1 advantages and disadvantages

Tip

Forest trust direction:

When configuring the forest trust, ensure it is a one-way trust where the production forest trusts the management forest. The management forest should have only an incoming trust. This prevents compromised production accounts from authenticating

against your backup infrastructure. Never configure a two-way trust between the management and production forests.

Option 2: Workgroup deployment

A workgroup deployment removes all Active Directory dependencies from the Veeam infrastructure. Each server is configured independently with local user accounts, local security policies, and local group memberships. This model is particularly well suited to smaller environments or organizations that require absolute isolation from any domain infrastructure.

The primary security advantage is simplicity of isolation. There are no domain trust relationships to exploit, no domain-cached credentials to extract, and no group policy inheritance to misconfigure. The backup infrastructure is entirely self-contained. However, this simplicity comes at the cost of management overhead that scales poorly. Every user account, password policy, audit configuration, and security setting must be configured individually on each server.

A significant technical limitation of the workgroup model is the loss of Kerberos authentication. All authentication between workgroup servers uses NTLM, which is considered less secure and is increasingly deprecated by Microsoft. Additionally, gMSAs require Active Directory and cannot be used in a workgroup environment, which limits options for automated credential rotation for guest processing operations.

Advantages	Disadvantages
Fast and straightforward to deploy	Significant management overhead in larger environments
Complete separation from production domain accounts	NTLM-only authentication (no Kerberos support)
No dependency on any domain infrastructure for availability.	gMSA accounts are unavailable for guest processing.
No additional domain controller infrastructure is required	Harder with a workgroup to demonstrate and maintain regulatory compliance
Minimal attack surface from a domain perspective	Insider threat response requires individual server-by-server lockouts

Technet24

Advantages	Disadvantages
	No centralized audit logging or Group Policy enforcement

Table 6.5: Option 2 advantages and disadvantages

Tip

Best fit:

A workgroup deployment is most practical for small environments with a limited number of Veeam infrastructure components (typically fewer than five servers) and a small team of backup administrators. For larger environments, the management overhead and compliance challenges generally outweigh the isolation benefits.

Option 3: Production domain (hardened)

Joining Veeam infrastructure to the production domain is the least secure option, but it is also the most operationally convenient and the deployment model commonly found in existing environments. If this approach is chosen, whether due to organizational constraints, resource limitations, or legacy architecture, it is essential to apply rigorous hardening measures to mitigate the elevated risk.

The core risk is that a breach of the production domain inherently compromises the backup infrastructure. Attackers who gain domain admin or equivalent privileges in the production forest can authenticate to the Veeam Backup & Replication server, enumerate all backup jobs and targets, and potentially destroy or encrypt backup data as part of a coordinated ransomware attack. This is precisely the scenario that a separate management forest or workgroup deployment is designed to prevent.

If production domain membership is unavoidable, the following hardening measures should be treated as mandatory rather than optional:

- **Enable MFA for all backup administrator accounts** using the built-in TOTP support in Veeam Backup & Replication v13
- **Place Veeam servers in a dedicated Organizational Unit (OU)** with a restricted group policy that blocks GPO inheritance from parent OUs and applies a hardened security baseline
- **Implement tiered administration** to ensure that production domain administrator credentials are never used to log in to Veeam infrastructure servers

- **Restrict remote access** to the Veeam Backup & Replication server using firewall rules, jump servers, or privileged access workstations (PAWs)
- **Monitor authentication events** on Veeam servers closely, forwarding security event logs to a SIEM that is outside the production domain's administrative reach
- **Deploy network segmentation** to restrict lateral movement between the production environment and the backup infrastructure subnet

Note

Critical risk:

Production domain membership means that any compromise of domain admin credentials in the production forest grants the attacker direct access to your backup infrastructure. This is the single most common vector exploited in ransomware attacks that specifically target backup systems. If you must use this model, treat every hardening measure as non-negotiable.

Securing service accounts

Service accounts used by Veeam Backup & Replication represent high-value targets for attackers because they typically possess elevated privileges across multiple systems. Veeam v13 requires service accounts for a range of operations, including backup job execution, guest OS processing, application-aware processing for SQL server, exchange, and Active Directory workloads, and communication with hypervisor management platforms.

The overarching principle for service account security is **least privilege combined with automated credential management.** Every account should have only the minimum permissions required for its specific function, and, wherever possible, the operating system should handle password rotation rather than relying on manual administrative processes.

Group managed service accounts (gMSA)

gMSAs are the preferred credential type for Veeam guest processing operations in Active Directory environments. A gMSA is a domain account whose password is automatically managed by the Windows operating system. Complex 240-character passwords are randomly generated and rotated every 30 days by default, eliminating the risk of brute force attacks, dictionary attacks, and stale credentials.

In Veeam Backup & Replication v13, gMSAs are supported for guest processing tasks, including application-aware processing for Microsoft Active Directory (domain controllers), Microsoft Exchange, Microsoft SQL Server, and Oracle 12c Release 2 and later. They cannot currently be used to run core Veeam services (such as the Veeam backup service itself) or for vCenter/

Technet24

hypervisor authentication, but they significantly reduce the credential attack surface for the guest interaction workflow.

Deploying gMSAs requires several prerequisites:

- **Active Directory domain services:** The key distribution service (KDS) must be enabled on at least one domain controller and a KDS root key must be generated before any gMSA passwords can be created.
- **Windows Server 2012 or later:** All machines that will use the gMSA (the Veeam backup server, guest interaction proxies, and target guest OS machines) must be running Windows Server 2012 or later.
- **Domain membership:** Both the guest interaction proxy and the target machine must have network access to the domain controllers and reside in the same domain as the gMSA to retrieve the managed password.
- **Local administrators group membership**: The gMSA must be a member of the local administrators group on each target machine. Domain administrator permissions are only required when backing up Active Directory domain controllers. Local administrator permissions are sufficient for all other supported applications.
- **Logon as a service:** The gMSA must be granted the "*Log on as a service*" permission on all machines where it will be used.

Note

gMSA creation example:

To create a gMSA on a domain controller, first generate a KDS root key (if one does not already exist): `Add-KdsRootKey -EffectiveImmediately`. Then create the account: `New-ADServiceAccount -Name "VBRgMSA" -DNSHostName "VBRgMSA.yourdomain.com" -PrincipalsAllowedToRetrieveManagedPassword "VeeamServers"`, where `VeeamServers` is an AD group containing the computer accounts that need to retrieve the gMSA password.

Standard service account hardening

For operations where gMSAs are not supported, such as running the Veeam backup service, authenticating to vCenter server, or configuring Linux-based repository connections, standard

Active Directory or local service accounts are required. These accounts demand careful attention to privilege management and credential hygiene:

- **Unique accounts per function:** Never reuse a single account across multiple roles. Create separate accounts for the Veeam backup service, hypervisor connections, repository access, and guest processing. This limits the blast radius if any single credential is compromised.
- **Deny interactive logon:** Configure service accounts with the "*Deny log on locally*" and "*Deny log on through Remote Desktop Services*" user rights. This prevents attackers from using stolen service account credentials for interactive sessions.
- **Enforce password rotation:** Implement a policy that rotates service account passwords at least every 90 days. Use a privileged access management (PAM) solution to automate this process where possible.
- **Remove domain admins membership:** No Veeam service account should be a member of the domain admins, enterprise admins, or schema admins groups. Grant only the specific permissions documented in the Veeam help center for each role.
- **Audit credential usage:** Enable and forward Windows security event log entries for logon (Event ID 4624), logon failure (Event ID 4625), and credential validation (Event ID 4776) on all servers where Veeam service accounts are used.

Multi-factor authentication (MFA)

Veeam Backup & Replication v13 includes built-in multi-factor authentication based on the time-based one-time password (TOTP) standard. When enabled, users must provide a six-digit code from a mobile authenticator application (such as Google Authenticator, Microsoft Authenticator, or Authy) in addition to their username and password when logging in to the Veeam console. Microsoft research indicates that accounts protected by MFA are 99.9% less likely to be compromised through credential-based attacks, making this one of the single most impactful security controls available.

Enabling MFA in the Veeam console

MFA is configured through the Veeam Backup & Replication console by a user with the Veeam Backup Administrator role. The process involves navigating to **Users and Roles**, opening the **Security tab**, ensuring that only individual user accounts (not groups) are listed, and selecting the **Enable multi-factor authentication (MFA)** checkbox. Once enabled, each user will be prompted to configure their authenticator application the next time they log in.

Technet24

Several important limitations and considerations apply to the console MFA implementation:

- MFA can only be enabled for individual user accounts, not for Active Directory groups. Any groups in the security list must be replaced with individual accounts before MFA can be activated.
- MFA is not supported for non-interactive connections used by backup infrastructure components, the PowerShell module, or automated scripts. These connections must use designated service accounts with MFA explicitly disabled.
- MFA is not natively available for Veeam backup enterprise manager. Enterprise Manager environments should use SAML-based authentication with a third-party identity provider that supports MFA, such as Duo Security, Okta, or Microsoft Entra ID.
- If a user loses access to their authenticator device, a backup administrator can reset their MFA enrollment through the security tab. The user will be prompted to re-enroll on their next login.
- After five consecutive failed MFA attempts, the account is temporarily locked. The user must wait at least one minute before attempting to log in again.

MFA for the Veeam software appliance host management UI

In v13 deployments using the VSA, there is a separate MFA enrollment for the host management Web UI (accessible on port 10443). This is independent of the backup console MFA and generates its own entry in the authenticator application. The security officer role on the appliance has MFA permanently enabled, and it cannot be disabled, reflecting the elevated privilege level of that role. Host administrator accounts can have MFA toggled on or off as needed.

It is important to understand that these are two distinct MFA configurations that must each be set up separately. Enabling MFA in the backup console does not automatically protect the Host Management UI, and vice versa. Both should be enabled in production deployments.

Note

Service account MFA exemption:

Accounts designated as service accounts (with the **This is a service account** checkbox selected in the Security settings) are exempt from MFA. This exemption is necessary for automated operations and PowerShell scripts, but it also means these accounts are less protected. Apply all other hardening measures rigorously to service accounts and monitor their usage closely for anomalous activity.

Network segmentation and DNS considerations

Regardless of which deployment model is selected, network segmentation is a critical complementary control. Veeam infrastructure components should be placed on a dedicated management network segment or VLAN, with firewall rules that restrict traffic to only the specific ports and protocols required for backup operations. This limits lateral movement if either the production network or the management network is compromised.

DNS resolution requires particular attention in separated environments. If the Veeam infrastructure uses a management forest or workgroup, it must maintain its own DNS infrastructure rather than relying on production DNS servers. Conditional forwarders or stub zones can be configured to enable name resolution across trust boundaries without creating a hard dependency on production DNS availability. In workgroup environments, a local hosts file or a lightweight DNS server dedicated to the management segment can provide resolution for Veeam infrastructure hostnames.

For environments where the VSA is deployed, the appliance's built-in network configuration tools allow the administrator to define DNS settings, NTP sources, and network interface bindings independently of any domain infrastructure, further reinforcing the principle of a self-contained management infrastructure.

The following table consolidates the key recommendations from this section into a prioritized checklist.

Priority	Recommendation
High	Deploy Veeam infrastructure in a separate management forest or workgroup isolated from the production Active Directory domain.
High	Enable multi-factor authentication for all interactive backup administrator accounts in both the Veeam console and the host management web UI.
High	Use group managed service accounts (gMSAs) for all supported guest processing operations to eliminate static passwords.
High	Implement network segmentation with firewall rules that restrict access to and from the Veeam management subnet.
High	Remove all Veeam service accounts from Domain Admins, Enterprise Admins, and other highly privileged groups.

Technet24

Priority	Recommendation
Medium	Configure one-way forest trust (production trust management) if cross-forest authentication is required.
Medium	Deploy dedicated DNS and NTP infrastructure for the management forest or workgroup segment
Medium	Enforce service account password rotation every 90 days or fewer using a PAM solution.
Medium	Forward all authentication event logs from Veeam servers to a SIEM outside the production domain.
Low	Deny interactive logon rights to all Veeam service accounts.
Low	Document all local accounts and security configurations in workgroup deployments for compliance and auditability.

Table 6.6: Priority checklist for security

Summary

This chapter has provided you with an understanding of using Active Directory or a workgroup in Veeam Backup & Replication v13. We discussed how to use Active Directory with Veeam Backup & Replication v13. Then we looked at best practices for Active Directory with Veeam Backup & Replication v13. Lastly, we examined how to set up and secure Active Directory for use with Veeam Backup & Replication v13. In the next chapter, we will cover the new High Availability with Veeam Backup & Replication v13 using VSA.

Further reading

- Veeam Backup & Replication – Domain or Workgroup Best Practice - `https://bp.veeam.com/security/Design-and-implementation/Hardening/Workgroup_or_Domain.html`
- Veeam Help Center – Multi-Factor Authentication - `https://helpcenter.veeam.com/docs/vbr/userguide/mfa.html`
- Veeam Help Center – Using Group Managed Service Accounts - `https://helpcenter.veeam.com/docs/vbr/userguide/using_gmsa.html`

- Veeam Help Center – Configuring Users (Host Management) - `https://helpcenter.veeam.com/docs/vbr/em/hmc_configure_users.html`
- Veeam Community Resource Hub – Security Functions & Features: MFA - `https://community.veeam.com/cybersecurity-170/security-functions-features-multi-factor-authentication-10379`
- Microsoft Learn – Group Managed Service Accounts Overview - `https://learn.microsoft.com/en-us/windows-server/identity/ad-ds/manage/group-managed-service-accounts/group-managed-service-accounts/group-managed-service-accounts-overview`

Technet24

7

Veeam High Availability (HA)

Veeam High Availability (HA) is part of the new **Veeam Software Appliance**, the newest release for v13. In this chapter, we will discuss what High Availability is, how to configure it, how failover works with HA, and how this can help your disaster recovery (DR) strategy to ensure resilience.

We're going to cover the following main topics:

- High Availability: Protecting your environment for disaster recovery
- High Availability cluster and how to configure it
- High Availability failover

Technical requirements

To follow along with this chapter, you should have access to the VSA ISO and two deployed Veeam Software Appliances.

- You must deploy two Veeam Software Appliances to configure High Availability.
- You must download the latest **ISO/OVA file** from `www.veeam.com`, which requires registration on the site and allows you to obtain a trial license. At the time of this writing, version `13.0.1.1071` is the current release.
- *Veeam Documentation* website to familiarize yourself with the official guidelines: `https://helpcenter.veeam.com/docs/vbr/userguide/high_availability_cluster.html?ver=13`.

Technet24

High Availability: Protecting your environment for disaster recovery

The backup server is the single most critical component in any Veeam Backup & Replication deployment. It orchestrates every backup, replication, and restore job across the entire infrastructure. If the backup server goes down, whether due to hardware failure, operating system corruption, or a site-level disaster, all data protection operations halt. Prior to version 13.0.1, recovering from a backup server outage meant deploying a fresh installation, restoring the configuration database from a backup, and re-establishing connectivity to every managed component. Depending on the size and complexity of the environment, this process could take hours.

With the introduction of HA clustering in Veeam Software Appliance (VSA) 13.0.1, Veeam has fundamentally changed the disaster recovery posture of the backup infrastructure itself. The HA feature deploys the backup server in an active/passive two-node cluster, with continuous database replication between nodes. When the primary node becomes unavailable, the secondary node can assume all operations with minimal interruption, eliminating the prolonged recovery windows that previously defined backup server failures.

Why HA matters for disaster recovery

Disaster recovery planning has traditionally focused on protecting production workloads such as virtual machines, databases, applications, and file servers. The backup server itself was often treated as a secondary concern, with organizations relying on configuration backups and documented rebuild procedures as their safety net. This approach introduces a dangerous gap: in a true disaster scenario, the very system responsible for executing recovery operations is itself unavailable.

The HA cluster addresses this gap by making the backup server a first-class participant in the disaster recovery strategy. Rather than recovering the recovery tool first, administrators can switchover or failover to the standby node and immediately begin restoring production workloads. This is particularly critical in ransomware incident response scenarios, where every minute of downtime directly impacts the blast radius of an attack and the organization's ability to contain and remediate the threat.

Note

Key insight:

The HA cluster does not protect your backup data—that remains the responsibility of your backup repositories, immutability settings, and offsite copies. What HA protects is

> the management plane which includes the configuration database, job definitions, credential stores, and the orchestration engine that makes recovery possible.

Overview of HA cluster architecture

The HA cluster is built on a two-node active/passive architecture. The primary node handles all backup, replication, and restore operations, while the secondary node maintains a continuously synchronized copy of the configuration database and stands ready to assume the primary role. This architecture leverages several key components working together to provide seamless failover capability.

Core components

Table 7.1 presents the core components in the HA cluster architecture:

Component	Role in the HA cluster
`veeamhasvc.service`	The HA service is responsible for cluster assembly, monitoring, failover, switchover, and disassembly operations. It runs on both nodes and coordinates all HA-related activities.
PostgreSQL instance	The configuration database that stores all job definitions, credentials, backup metadata, and infrastructure mapping. Operates in read-write mode on the primary node and read-only mode on the secondary node to prevent split-brain scenarios.
WAL file replication	Write-Ahead Log (WAL) files capture every database transaction on the primary node and are continuously replicated to the secondary node. If the secondary node is temporarily unavailable, WAL files are stored in `/var/lib/veeamdata/pgsql/pg_wal` on the primary until synchronization resumes.
Virtual cluster IP	A floating IP address assigned to whichever node is currently serving as primary. During switchover or failover, this IP migrates to the new primary node, ensuring that the Veeam Backup & Replication console and all infrastructure components continue to reach the active backup server without reconfiguration.

Technet24

Component	Role in the HA cluster
Cluster DNS Name	A DNS name that is mapped to the cluster endpoint. This provides an alternative to the virtual IP for console connections and can simplify certificate management and firewall rules in enterprise environments.
Console Local Cache	The Veeam Backup & Replication console maintains a local cache of HA cluster information, including both node IP addresses and their current roles. This cache enables the console to detect the cluster state and is essential for initiating switchover or failover operations while preventing accidental failover when the primary is still running.

Table 7.1: HA cluster core components

Data synchronization and replication

The integrity of the HA cluster depends entirely on keeping the secondary node's database synchronized with the primary. Veeam leverages PostgreSQL's native replication capabilities to achieve this, running a continuous HA maintenance job that performs the following operations:

1. **Cluster state verification**: The PostgreSQL database service checks the HA cluster state and verifies that both nodes are available and communicating.
2. **WAL log capture**: The maintenance job saves WAL files on the primary node, capturing every database transaction since the last synchronization cycle.
3. **Data transfer**: The maintenance job transmits a snapshot of the Veeam configuration database, WAL files, client certificates, configuration files for registry-based settings, and synchronization inclusion/exclusion rules to the secondary node.
4. **WAL application**: The maintenance job overwrites existing WAL logs on the secondary node with the latest files, bringing the secondary database to an identical state.

Note

Important:

If the secondary node becomes unavailable for more than 10 minutes, Veeam Backup & Replication triggers notification mechanisms including an alert in the notification bar, email notifications (if configured), and SNMP traps. When the secondary node is restored, a follow-up notification confirms the cluster has returned to a healthy state.

Split-brain prevention

One of the most dangerous scenarios in any clustered system is a split-brain condition, where both nodes simultaneously believe they are the primary and begin accepting writes independently. This leads to data divergence and can result in irrecoverable corruption of the configuration database.

Veeam prevents this through a deliberate architectural decision: the primary node always operates in **read-write mode**, while the secondary node is enforced to **read-only mode**. The secondary node runs no jobs and no services beyond those required for database replication and communication with the primary. This means there is no scenario in which both nodes can simultaneously modify the configuration database.

Additionally, the Veeam Backup & Replication console's local cache plays a protective role. Before initiating a failover, the console verifies the state of both nodes using its cached information, helping to prevent an accidental failover to the secondary node when the primary is still operational. This is a critical safeguard that prevents operator error from creating a split-brain condition.

Disaster recovery design patterns with HA

The HA cluster's ability to tolerate high-latency synchronization opens several powerful disaster recovery design patterns. The secondary node does not need to be co-located with the primary, enabling geographically distributed deployments that protect against site-level failures.

Pattern 1: Same-site HA for hardware and OS failures

Deploy both VSA nodes within the same data center, ideally on different physical hosts and storage subsystems. This pattern provides protection against server hardware failures, storage controller issues, and operating system corruption. It is the simplest deployment model and offers the fastest failover times due to low-latency synchronization between nodes.

Pattern 2: Cross-site HA for site-level disasters

Place the secondary VSA node at a remote data center or disaster recovery site. Because the HA cluster supports synchronization across high-latency WAN links, this pattern provides genuine site-level disaster recovery for the backup management plane. In the event of a complete site failure, administrators can failover to the remote node and begin executing restore operations against offsite backup copies, replicas, or cloud-tiered data.

Technet24

Pattern 3: Combined HA with Veeam ONE automated failover

Integrate Veeam ONE as a cluster witness monitoring the HA cluster state. Configure remediation actions to automatically trigger failover when the primary becomes unreachable. This pattern delivers hands-off disaster recovery for the backup server, minimizing the mean time to recovery (MTTR) even when administrators are not immediately available. It is particularly well-suited for managed service providers and organizations with limited after-hours staff.

Alternatives when HA cannot be deployed

Not every environment will meet the prerequisites for HA clustering. Organizations that cannot deploy HA due to licensing constraints (editions below Premium), platform limitations (Windows-based installations), or infrastructure restrictions should consider the following alternative approaches to protecting the backup server:

- **Configuration backup with standby server**: Maintain regular encrypted configuration backups stored on an external repository. Deploy a standby backup server (or document the deployment procedure) that can be brought online and have the configuration restored in an emergency. This is the traditional approach and remains viable, though recovery times are measured in hours rather than minutes.
- **VM replication of the backup server**: If the backup server runs as a virtual machine, use Veeam Backup & Replication to replicate it to another hypervisor host. In a disaster, manually start the replica VM. This approach provides faster recovery than configuration restore but requires the hypervisor infrastructure to be available.
- **Hypervisor-level HA**: Leverage VMware vSphere HA or similar hypervisor clustering features to automatically restart the backup server VM on another host in the event of an ESXi host failure. This protects against host hardware failures but does not protect against VM-level or guest OS issues.

HA clustering in VSA 13.0.1 represents a maturation of Veeam's approach to backup infrastructure resilience. By treating the backup server as a critical system deserving its own disaster recovery plan rather than an afterthought protected only by configuration backups, organizations can dramatically reduce the time required to resume data protection operations after an incident.

The combination of continuous PostgreSQL replication, virtual IP failover, split-brain prevention, centralized updates, and the option for Veeam ONE–automated failover creates a robust, enterprise-grade HA solution that integrates naturally with existing Veeam infrastructure. For organizations running the Premium edition of Veeam Data Platform on the Linux-based Software Appliance, enabling HA should be considered a foundational best practice for any production deployment.

High Availability (HA) cluster

One of the most significant features introduced in Veeam Backup & Replication v13.0.1 running on the Veeam Software Appliance is the **High Availability (HA) cluster**. This capability addresses a long-standing gap in the Veeam architecture: the backup server itself was historically a single point of failure. While Veeam has always excelled at protecting workloads across virtual, physical, and cloud environments, the management server coordinating all those operations had no native redundancy mechanism. If the backup server went down, backup and replication jobs would halt until it was restored, typically by deploying a new server and importing a configuration backup.

The HA cluster feature changes this paradigm entirely. By clustering two VSA nodes into an active/passive pair, organizations gain continuous availability of the backup server with minimal downtime during planned maintenance or unexpected failures.

Understanding High Availability

The Veeam HA cluster is a two-node clustering solution designed to ensure continuous availability and minimize downtime of the backup server. It works by pairing two Linux-based VSA instances into an active/passive configuration where one node (the primary) handles all backup and replication operations while the second node (the secondary) stands by in a synchronized state, ready to assume the primary role when needed.

The core principle is straightforward: the configuration database on the primary node is continuously replicated to the secondary node using PostgreSQL streaming replication. If the primary node becomes unavailable, the secondary node can be promoted to take over operations with all job configurations, credentials, encryption keys, and scheduling information intact. There is no need to restore a configuration backup, re-enter credentials, or reconfigure jobs. The backup infrastructure simply continues to operate.

Figure 7.1 depicts the HA cluster components:

Technet24

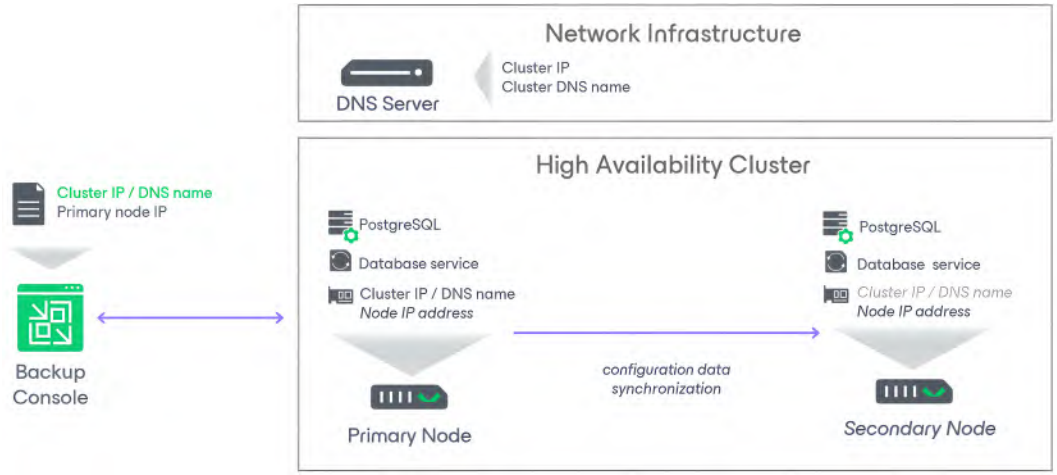


Figure 7.1: High-Availability cluster components

Let's look at the key components that make up the HA cluster:

- **Cluster IP address (Virtual IP):** A dedicated IP address assigned to the HA cluster that always points to the currently active primary node. Administrators connect to this address rather than to individual node IPs. During switchover or failover, the cluster IP is automatically reassigned to the new primary node, ensuring transparent connectivity.
- **Cluster DNS name:** A DNS hostname associated with the cluster that can be used instead of the virtual IP address. After assembling the cluster, all console connections should use the cluster DNS name or IP rather than individual node addresses this way if there is an automatic failover or a manually required failover you always connect to the Primary node.
- **Primary node:** The active node that runs all Veeam Backup & Replication services and handles backup, replication, and restore operations. The primary node maintains the writable copy of the PostgreSQL configuration database.
- **Secondary node:** The standby node that maintains a continuously synchronized read-only copy of the configuration database. All Veeam services except those required for database replication and communication with the primary are stopped on this node.

Now that we have looked at what the key components of the HA cluster are we can dive into how the HA cluster works.

How the HA cluster works

Now that you understand the components that make up the HA cluster, we can now look at the required steps to enable the feature in your environment.

Cluster assembly

When the HA cluster is assembled, Veeam Backup & Replication performs a series of operations on both nodes to establish the replication relationship:

- On both nodes, Veeam starts the `veeamhasvc.service` (the HA service responsible for cluster assembly, failover, switchover, and disassembly operations). It also opens the necessary ports for PostgreSQL database communication between nodes.
- On the primary node, Veeam assigns the virtual cluster IP address and begins generating WAL (Write-Ahead Log) files for database replication.
- On the secondary node, Veeam removes existing data from the PostgreSQL database, sets the PostgreSQL database to read-only mode, and stops all jobs and services except those necessary for database replication. This is done to synchronize the database from the primary node to the secondary node.

Veeam Backup & Replication then creates a local cache containing information about the cluster configuration and initiates continuous synchronization of the configuration databases between the two nodes.

Data synchronization

The synchronization process leverages PostgreSQL's native streaming replication capabilities. The HA maintenance job continuously performs the following operations:

- **PostgreSQL service monitoring**: The database service constantly checks the PostgreSQL configuration and verifies that both HA nodes are available.
- **WAL log management**: The maintenance job saves Write-Ahead Log files on the primary node and transmits them to the secondary node.
- **Configuration snapshot transfer**: A snapshot of the Veeam Backup & Replication configuration database is sent to the secondary node.
- **Credential and certificate sync**: Client certificates and credential data are synchronized to ensure the secondary node has everything needed to assume the primary role.
- **Configuration file replication**: Registry-based settings and HA cluster configuration files that specify inclusion/exclusion rules for synchronization are transmitted.

Technet24

Note

Important:

Enabling HA synchronizes all configuration database content including credentials and encryption secrets to the secondary cluster node. Because of this, assembling an HA cluster requires approval from the security officer (if one was configured during the initial VSA setup). This is a deliberate security gate to prevent unauthorized replication of sensitive data. If there was no security officer role defined during initial VSA setup the enablement of HA is automatically turned on. You can reference *chapters 1 and 2* for the security officer role and setup.

WAL file storage and resilience

If the secondary node becomes temporarily unavailable, the primary node does not halt operations. Instead, it continues to store the replication delta as WAL files in the `/var/lib/veeamdata/pgsql/pg_wal` directory. When the secondary node comes back online, the accumulated WAL files are replayed to bring it back into sync. If the secondary node is unavailable for more than 10 minutes, Veeam Backup & Replication displays a notification in the console to alert the administrator.

Infrastructure requirements

Before assembling an HA cluster, several infrastructure components must be in place. Careful planning at this stage prevents configuration issues during cluster assembly.

Licensing

The HA cluster feature requires the **Veeam Data Platform — Premium Edition** license. This is the package that includes Veeam Backup & Replication, Veeam ONE, and Veeam Recovery Orchestrator. The license only needs to be installed on the primary node. It is replicated to the secondary node as part of the database synchronization.

Server requirements

Table 7.2 consolidates the server requirements to assemble an HA cluster:

Requirement	Details
Operating system	Veeam Software Appliance (Linux-based, JeOS)
Primary node	May already have Veeam Backup & Replication installed and operational

Requirement	Details
Secondary node	Must be a freshly deployed Veeam Software Appliance with no existing configuration
Network subnet	Both nodes must reside within the same IP subnet
IP addressing	Three static IPs required: one for each node, plus one for the cluster virtual IP
DNS	Forward and reverse lookup zones with records for cluster endpoint, primary node, and secondary node

Table 7.2: Server requirements for HA cluster nodes

Network and DNS configuration

Proper DNS configuration is critical for HA cluster operation. You must create forward and reverse lookup records for three endpoints:

- **Cluster DNS name**: resolves to the virtual cluster IP address
- **Primary node hostname**: resolves to the primary node's static IP
- **Secondary node hostname**: resolves to the secondary node's static IP

Both nodes must be able to resolve each other by hostname, and the cluster virtual IP must not be assigned to any other device on the network. The virtual IP will be dynamically assigned to whichever node currently holds the primary role.

Tip

Before starting the cluster assembly wizard, verify DNS resolution from both nodes. Use `nslookup` or `dig` to confirm that forward and reverse lookups resolve correctly for both node hostnames and the planned cluster DNS name. DNS misconfiguration is one of the most common causes of HA cluster assembly failures.

HA cluster internal components

Once assembled, the HA cluster relies on several built-in services and components running on both nodes. Understanding these components is essential for troubleshooting:

Technet24

Component	Purpose
`veeamhasvc.service`	The core HA service. Manages cluster assembly, disassembly, failover, and switchover operations.
PostgreSQL instance	Contains the Veeam configuration database. Runs in read-write mode on primary, read-only on secondary. Uses streaming replication for synchronization.
WAL files storage `/var/lib/veeamdata/pgsql/pg_wal`	Stores Write-Ahead Log deltas on the primary node when the secondary is temporarily unavailable, enabling catch-up synchronization.
VBR console cache	Stores cluster configuration locally on the console machine at `C:\Users\<USER>\AppData\Local\Veeam_Software_Group_GmbH`. Used for connection management and node status tracking.

Table 7.3: HA cluster internal components

Configuring the HA cluster

The HA cluster configuration process involves two major phases: enabling the HA feature on both nodes via the Host Management Console, and then assembling the cluster through the Veeam Backup & Replication Console.

Phase 1: Enabling HA on both nodes

Before the cluster can be assembled, the HA capability must be enabled on each Veeam Software Appliance independently through the host management console (cockpit web interface on port 10443).

Follow these steps to enable HA capability:

1. Access the host management console of the first VSA node by navigating to `https://<node-ip>:10443` in a web browser.

2. Navigate to the **Integrations** | **Backup Infrastructure** and locate the **High Availability** heading.
3. Click the **Submit Request** link to enable HA. Since this operation exposes configuration database content (including credentials) to a second node, the security officer must approve the request. (In my lab, I did not enable the Veeam security officer account, so did not require the approval request.)

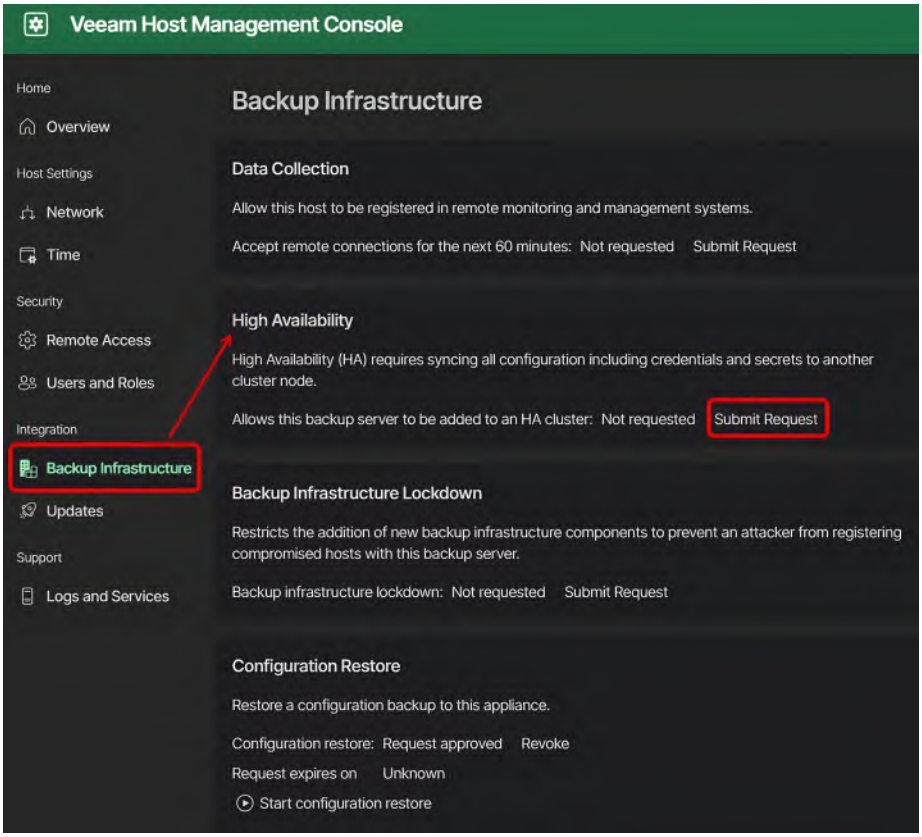


Figure 7.2: Enabling HA on the management Web UI

4. Approve the request by logging in as the Security Officer (if enabled) and confirm the HA enablement.

Technet24

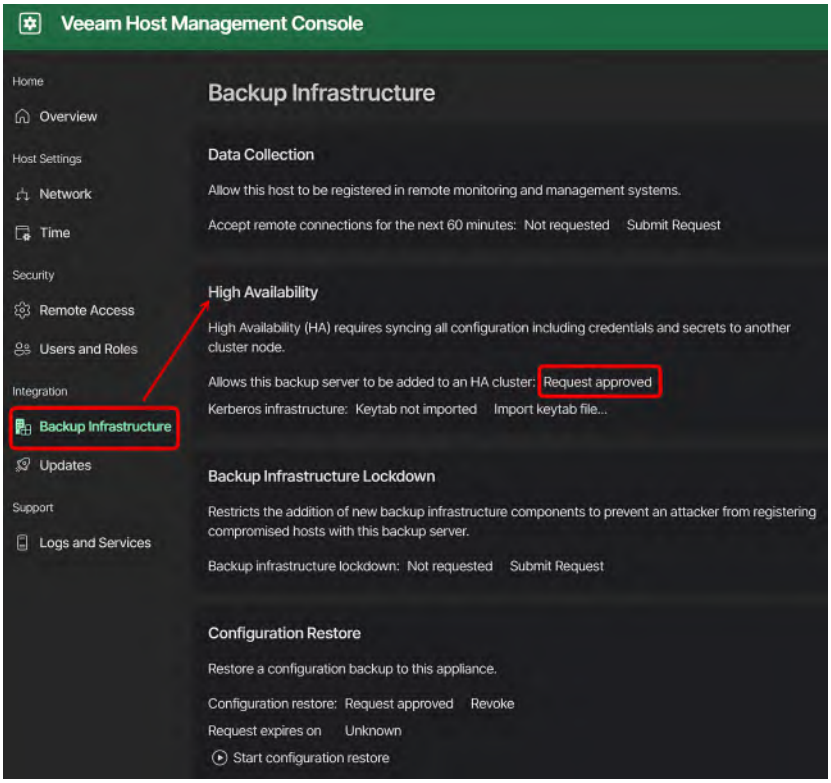


Figure 7.3: High Availability request showing approved

5. Repeat the process on the second VSA node. Both nodes must have HA enabled before cluster assembly can proceed.

Note

Important:

If Multi-Factor Authentication (MFA) is enabled on the host management console (which is recommended), ensure that OTP codes are configured for both the administrator and Security Officer accounts on both nodes before proceeding. Time synchronization between nodes must also be accurate for OTP validation to succeed.

Phase 2: Assembling the cluster via the VBR console

Once HA has been enabled and approved on both nodes, the cluster is assembled using the New High Availability Cluster wizard in the Veeam Backup and Replication Console (the Windows-based remote management console).

Step 1: Launching the wizard

Follow these steps to launch the wizard:

1. Connect the Veeam Backup & Replication Console to the machine that will serve as the primary node.
2. Open the **Backup Infrastructure** view, select **Managed Servers** in the inventory pane, then select the Linux host and click **Create HA cluster** on the ribbon (or right-click the host and select **Create HA cluster**).

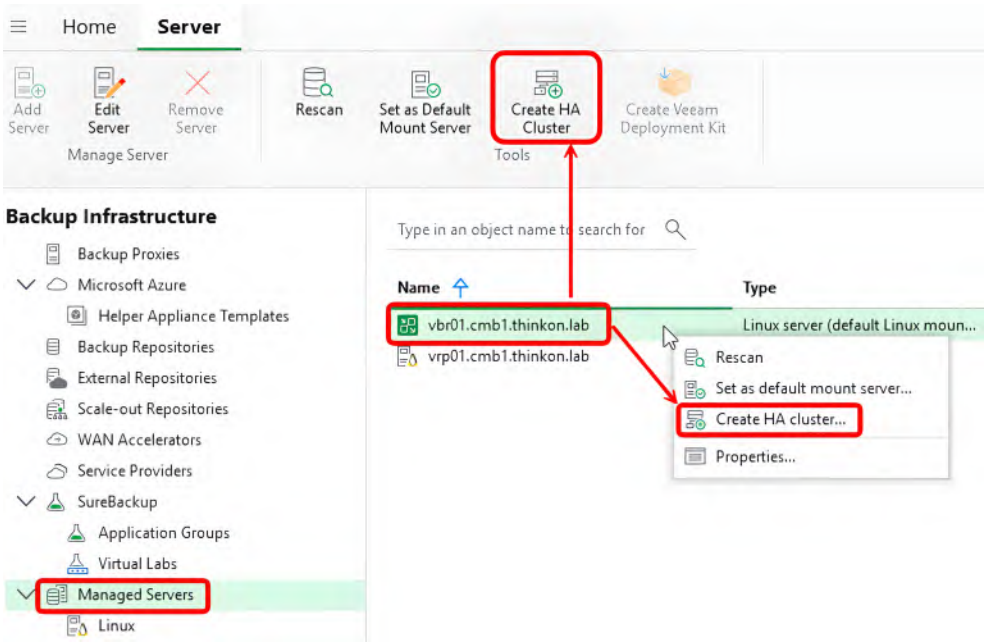


Figure 7.4: Create HA cluster options in toolbar and right-click menu

Step 2: Specifying cluster settings

Follow these steps to specify cluster settings:

1. Enter the **Cluster DNS name** and the **Virtual IP address**. The DNS name should match the DNS record you created during the preparation phase.
2. Click **Next** to proceed.

Technet24

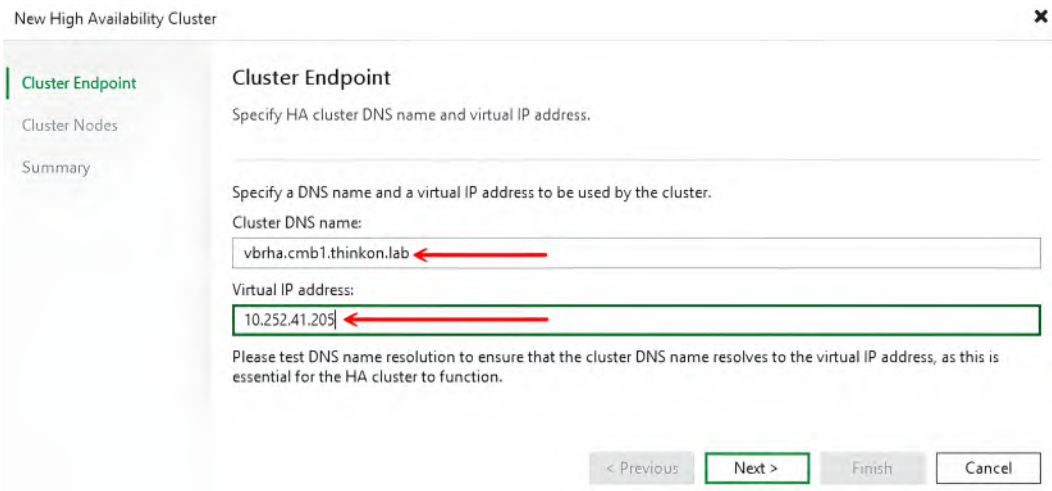


Figure 7.5: Cluster FQDN in DNS and virtual IP Address

Step 3: Specifying cluster node settings

These steps will help you navigate and specify cluster node settings:

1. Select the IP address of the first node (the current host), then enter the IP address or hostname of the second VSA node along with credentials for authentication.

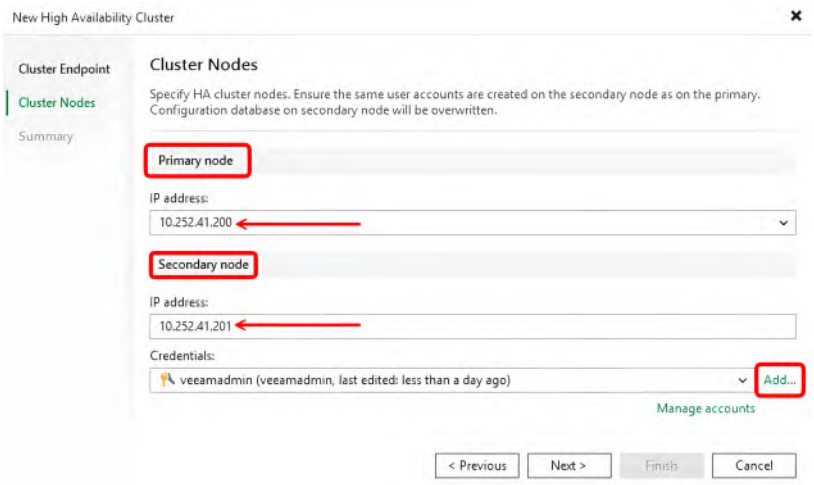


Figure 7.6: Cluster node information for primary and secondary

2. Click **Next** to verify connectivity and accept the **Certificate Security Alert** by clicking **Continue**

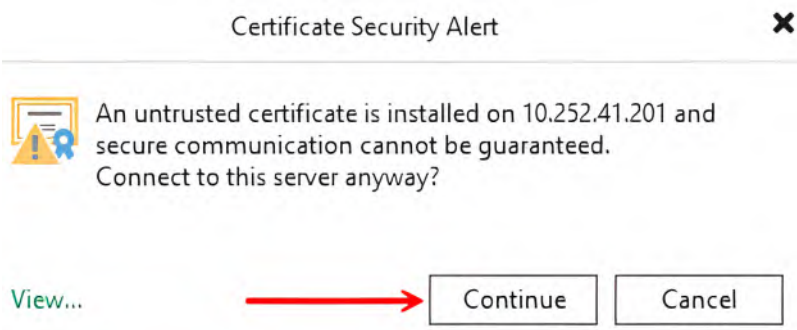


Figure 7.7: Certificate warning – click Continue to connect

Step 4: Finishing the wizard

With this, you'll complete the HA cluster configuration:

1. Review the summary and click **Finish** to begin cluster assembly. Veeam Backup & Replication will configure both nodes, establish database replication, and assign the virtual IP. This process typically takes 10–15 minutes depending on the environment.

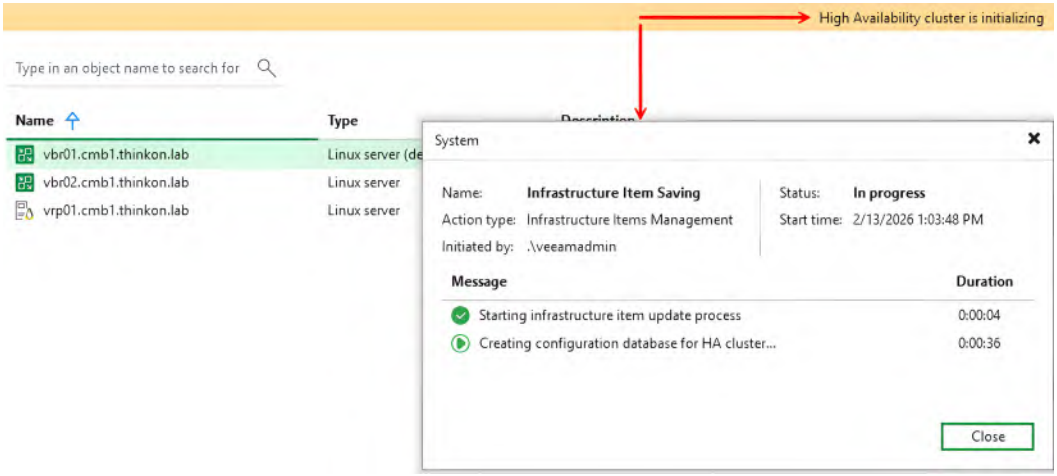


Figure 7.8: High Availability cluster is initializing

Once assembly completes, all status indicators in the console should show green. From this point forward, you must connect to the backup server exclusively using the **cluster DNS name** or **cluster IP address**—not the individual node addresses.

Technet24

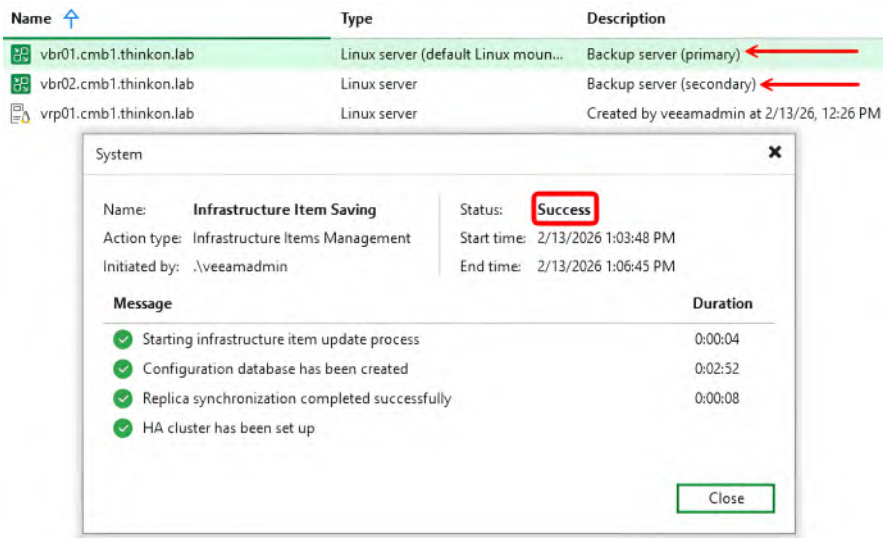


Figure 7.9: High Availability cluster configuration completed

Tip

There is a known issue in version 13.0.1 where HA cluster creation can fail if a capacity tier is configured or if the local default repository is still present. If you encounter errors during assembly, try removing the default local repository and pointing the configuration backup to an external repository before retrying the wizard.

Switchover and failover operations

The HA cluster supports two methods for transitioning the primary role from one node to another: **switchover** (planned) and **failover** (unplanned).

Switchover (planned)

A switchover is a graceful, administrator-initiated transition of the primary role to the secondary node. This is typically used for planned maintenance activities such as patching, hardware upgrades, or testing the failover capability.

During a switchover, the following actions take place:

- All running jobs are stopped gracefully
- The configuration database on the secondary is brought fully up to date
- The virtual cluster IP is reassigned from the current primary to the secondary node

- The former secondary assumes the primary role and starts all Veeam services
- The former primary transitions to the secondary (standby) role

To initiate a switchover, follow these steps:

1. Open the **Backup Infrastructure** view in the console, right-click one of the HA cluster nodes, and select **Switchover to Another Node**. You can also click the toolbar button labeled the same.
2. You will be prompted to disconnect from the console.
3. After the process completes (typically 10–15 minutes), reconnect using the cluster IP or DNS name. The console will automatically connect to the new primary.

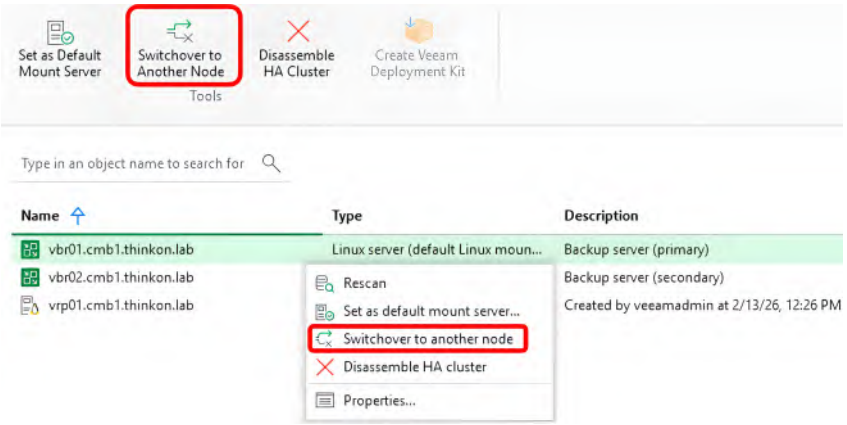


Figure 7.10: Planned failover/switchover to secondary node

The above option as noted starts the failover process from the current primary node and once the failover is completed, you will notice that the node which was the **"secondary"** previously is now the **"primary"** node. This changes the previous **"primary"** node to now be the **"secondary"** node in the cluster. When you select the option via right-click or the toolbar button, you do receive a warning about running jobs and need to click **Switchover** to begin.

Technet24

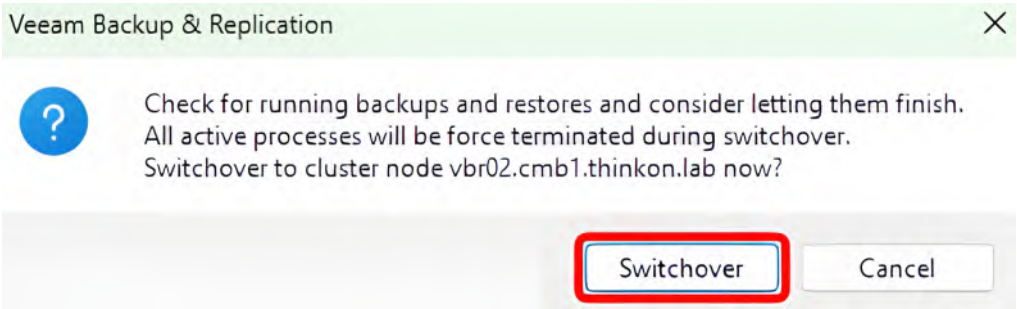


Figure 7.11: Switchover warning about backup and restore jobs

Below is the result of the switchover, with the nodes switching roles.

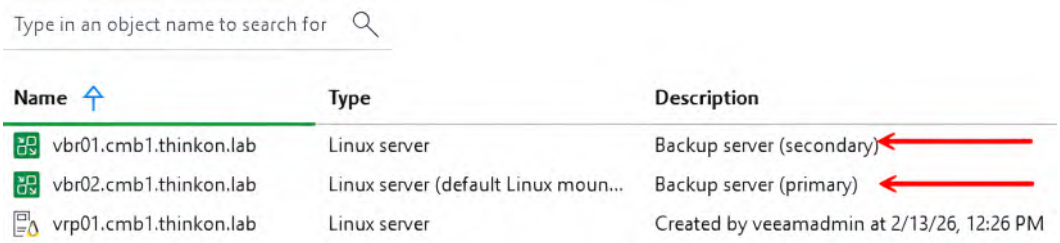


Figure 7.12: Switchover completed with new primary node

Failover (unplanned)

A failover is performed when the primary node has become unexpectedly unavailable due to hardware failure, network outage, or operating system crash. Unlike switchover, failover is a recovery operation where some data loss is possible if the most recent WAL files have not yet been replicated to the secondary.

During a failover, the following actions take place:

- The administrator manually initiates the failover from the console (automatic failover is not built-in by default.)
- The secondary node promotes its PostgreSQL database from read-only to read-write.
- The cluster virtual IP is assigned to the secondary node.
- Veeam services start on the promoted node.
- The node assumes the primary role with all synchronized configuration data.

Note

As of v13.0.1, failover must be manually initiated by an authorized administrator. There is no built-in automatic failover trigger. However, Veeam ONE can be configured to

monitor cluster health and, through policy-driven automation, can be used to trigger failover operations. Additionally, the availability of PowerShell cmdlets enables organizations to build custom witness-node- or heartbeat-based automation logic.

Switchover vs. failover

Table 7.4 draws the comparison between switchover and failover operations:

Aspect	Switchover	Failover
Trigger	Administrator-initiated (planned)	Administrator-initiated (reactive to failure)
Data loss risk	None — full synchronization before transition	Minimal — limited to unreplicated WAL files
Job handling	Jobs stopped gracefully	Jobs interrupted by a primary failure
Typical duration	10–15 minutes	5–10 minutes
Use case	Maintenance, patching, testing	Hardware failure, OS crash, network outage
Primary node status	Transitions to a secondary role	Remains offline until recovered

Table 7.4: Switchover versus failover comparison

Post-failover recovery

After a failover event, the original primary node must be recovered and reintegrated into the cluster. The procedure depends on the nature of the failure.

If the original primary can be brought back online, it will rejoin the cluster as the secondary node and begin receiving database replication from the new primary. You can then perform a switchover to return the original node to the primary role, if desired.

If the original primary cannot be recovered, deploy a fresh Veeam Software Appliance, enable HA on it, and use the cluster management interface to add it as the new secondary node. The configuration database will be replicated from the current primary.

Technet24

Managing the HA cluster

Let's now look at the best ways to manage the HA cluster once it is created and running.

Monitoring cluster status

HA cluster session data including creation, synchronization, switchover, and failover events can be viewed directly in the Veeam Backup & Replication console. Navigate to the **Home** view, select **Last 24 Hours** or **Running** in the inventory pane, and locate the **HA cluster Synchronization** session in the working area. Double-clicking the session opens detailed statistics and event logs.

Rescanning the cluster

If the cluster state appears inconsistent or after recovering a failed node, use the **Rescan** option (right-click the **Backup server (primary)** in **Backup Infrastructure**) to force a refresh of the cluster configuration and synchronization status.

Disassembling the cluster

The HA cluster can be disassembled through the **Backup Infrastructure** view if you no longer need high availability. This operation removes the cluster configuration and returns both nodes to standalone operation. The primary node retains the full configuration database, while the secondary node's database is cleared.

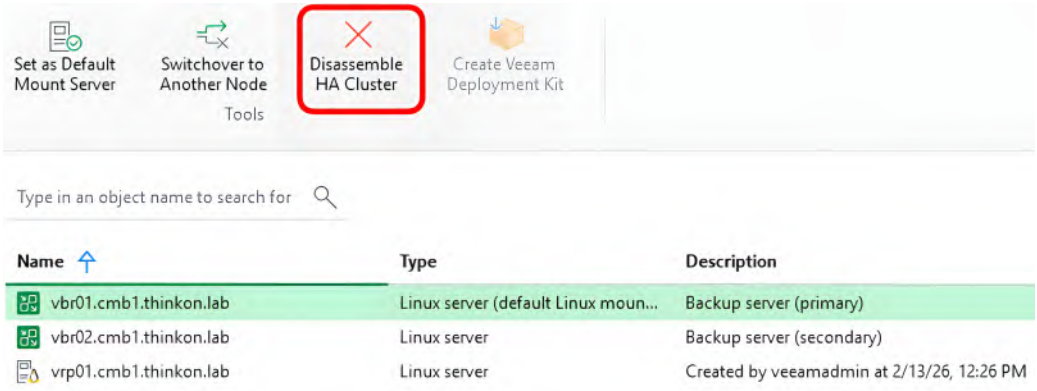


Figure 7.13: Disassemble HA cluster button

Automatic updates in an HA cluster

One of the significant operational advantages of the HA cluster is its ability to handle updates for both nodes in a coordinated manner. The following actions take place when an update is available from Veeam:

1. Veeam Backup & Replication updates the primary node first.

2. It identifies which updates have been installed on the primary.
3. It queries the secondary node and compares installed updates.
4. It creates a list of updates needed on the secondary.
5. It sends commands to the secondary node to install the required updates.

This coordinated update mechanism eliminates the risk of version mismatch between nodes and reduces the administrative overhead of maintaining the cluster.

Benefits of using the HA cluster

Implementing the HA cluster provides several significant advantages for enterprise backup environments. *Table 7.5* presents a consolidated list of its benefits:

Benefit	Description
Eliminates single point of failure	The backup server is no longer a single point of failure. If the primary node goes down, the secondary can take over with all configuration data intact.
Fast recovery	Recovery from a backup server failure takes minutes rather than the hours typically required to deploy a new server and restore a configuration backup.
No configuration backup dependency	During a failover, there is no need to import configuration backups, recreate jobs, import backup files, or re-enter credentials. Everything is already synchronized.
Zero maintenance downtime	Planned maintenance can be performed without backup window interruption by performing a switchover, patching the node, and switching back.
Automated dual-node updates	Software updates are coordinated across both nodes automatically, eliminating version mismatch risks and reducing administrative effort.
Credential and secret synchronization	All credentials, encryption keys, certificates, and secrets stored in the configuration database are continuously replicated, ensuring complete operational readiness on both nodes.
Transparent connectivity	The virtual cluster IP and DNS name provide a single access point that follows the active primary node, making the failover process transparent to management consoles and automation scripts.

Technet24

Benefit	Description
Geographic distribution support	The HA cluster supports replication across relatively high network latencies, enabling node placement across different physical locations within the same subnet for added resilience.

Table 7.5: HA cluster benefits

Considerations and limitations

While the HA cluster provides critical backup server resilience, it is important to understand what it does and does not protect:

- **HA is not a backup:** The HA cluster synchronizes the configuration database between nodes in near-real-time. If the configuration becomes corrupted on the primary (for example, due to a software bug or accidental misconfiguration), that corruption will be replicated to the secondary. Configuration backups remain essential as a separate recovery mechanism.
- **HA does not protect repositories:** The cluster provides resilience for the backup server management plane only. The backup data stored in repositories is not replicated between HA nodes. Repository-level protection should be addressed through separate mechanisms such as SOBR with capacity/archive tiers, hardened repositories, or storage-level replication.
- **Same subnet requirement:** Both cluster nodes must reside within the same IP subnet for the virtual IP mechanism to function correctly.
- **Linux-only (VSA):** The HA cluster is available exclusively for Veeam Software Appliance deployments on Linux. Windows-based Veeam Backup & Replication installations do not support this feature.
- **Manual failover by default:** Failover is not automatic out of the box. An authorized administrator must initiate the failover operation. Custom automation can be built using PowerShell cmdlets and external monitoring.
- **Premium licensing required:** The feature requires *Veeam Data Platform — Premium Edition*, which may not be available in all deployment scenarios.
- **Fresh secondary node required:** The secondary node must be a newly deployed Veeam Software Appliance with no pre-existing Veeam configuration. You cannot use an existing deployed Veeam Software Appliance with this configuration. If not a newly deployed Veeam Software Appliance, the HA cluster configuration will not work correctly or you might have issues.

The HA cluster in Veeam Backup & Replication v13.0.1 represents a major step forward in backup infrastructure resilience. By pairing two Veeam Software Appliance nodes into an active/passive cluster with continuous PostgreSQL database replication, organizations can eliminate the backup server as a single point of failure and achieve rapid recovery from both planned and unplanned outages.

The configuration process is wizard-driven and accessible, requiring two primary phases: enabling HA through the host management console on both nodes (with security officer approval), and then assembling the cluster through the VBR Console. Once operational, the cluster handles database synchronization, credential replication, and software updates automatically, significantly reducing the administrative overhead of maintaining a resilient backup infrastructure.

While the feature has certain limitations (notably the manual failover requirement, same-subnet constraint, and premium licensing dependency,) it provides compelling value for enterprise environments where backup server availability directly impacts recovery point and recovery time objectives.

High Availability failover

HA in Veeam Backup & Replication v13.0.1 ensures business continuity by maintaining a standby configuration database node that can assume the role of the active backup server in the event of a hardware failure, software crash, or planned maintenance window. In this section, we will examine the failover architecture, the mechanisms that govern automatic and manual failover, and the operational procedures required to manage HA transitions in production environments.

Understanding HA failover architecture

The Veeam HA cluster operates on a leader–follower model, where a single **active node** processes all backup, replication, and restore operations, while one or more **passive nodes** maintain synchronized copies of the configuration database and stand ready to assume the active role. The underlying technology stack relies on three critical components that work in concert to deliver seamless failover capability:

- **Patroni**: Manages PostgreSQL replication and orchestrates leader election across the cluster nodes. Patroni continuously monitors the health of the active database instance and triggers promotion of a standby replica when a failure is detected.

Technet24

- **etcd**: Provides the distributed consensus store that Patroni relies on for leader election. The etcd cluster maintains a consistent view of which node currently holds the leader lock, ensuring that only one node can be promoted to active status at any given time.
- **keepalived**: Manages the virtual IP (VIP) address that clients and managed components use to connect to the backup server. When Patroni promotes a new leader, keepalived migrates the VIP to the newly active node, ensuring transparent connectivity for all Veeam services.

Together, these components form a tightly integrated failover pipeline: etcd provides the consensus layer, Patroni manages the database promotion logic, and keepalived handles the network-level redirection. Understanding how these three services interact is essential for troubleshooting failover events and for planning maintenance operations that minimize disruption.

Failover event sequence

When the active node becomes unavailable, the HA cluster executes a deterministic sequence of operations to restore service. The following table outlines each phase of the automatic failover process:

Phase	Component	Action
1	Patroni	Detects loss of heartbeat from the active PostgreSQL instance. The failure detection timeout is governed by the `ttl` and `loop_wait` parameters in the Patroni configuration.
2	etcd	Releases the leader lock held by the failed node after the TTL expires. The remaining nodes are now eligible to compete for the leader lock.
3	Patroni	The standby node's Patroni instance acquires the leader lock from etcd and promotes its local PostgreSQL replica to read-write mode.
4	keepalived	Detects the Patroni leadership change and migrates the Virtual IP address to the newly promoted node via a gratuitous ARP announcement.

Phase	Component	Action
5	Veeam services	The Veeam Backup Service on the new active node detects the database promotion and begins accepting client connections. Running jobs that were in progress on the failed node are automatically restarted or resumed, depending on the job type and checkpoint availability.

Table 7.6: Failover components and actions taken

Note

The total failover time depends on several factors, including the Patroni TTL configuration, network latency between nodes, and the size of any write-ahead Log (WAL) segments that must be replayed on the promoted standby. In a well-configured two-node VSA cluster on a low-latency network, typical automatic failover completes within 30–60 seconds.

Automatic vs. manual failover

Veeam HA supports two distinct failover modes, each suited to different operational scenarios. Choosing the appropriate mode depends on the organization's tolerance for brief service interruptions versus the need for controlled transitions during planned maintenance.

Automatic failover

Automatic failover is the default behavior of the HA cluster. When the active node fails unexpectedly, whether due to a hardware fault, operating system crash, or network partition, the cluster autonomously promotes the standby node without operator intervention. This mode is designed for unplanned outages where speed of recovery takes precedence over administrative oversight.

Key characteristics of automatic failover include the following:

- No manual intervention required; Patroni and etcd handle the entire promotion sequence.
- The Virtual IP migrates automatically, so connected clients and managed agents reconnect transparently.

Technet24

- Running backup and replication jobs that support checkpointing will resume from their last checkpoint on the new active node.
- Jobs that do not support checkpointing (such as certain restore operations) will need to be restarted manually after failover completes.

Note

WARNING!

In a two-node HA cluster, if both nodes experience simultaneous failures (for example, due to a shared storage subsystem failure or a common network switch outage), automatic failover cannot recover the cluster. Always ensure that HA nodes are deployed on physically separate infrastructure with independent power, network, and storage paths.

Manual (planned) failover

Manual failover, also referred to as a *switchover*, is an operator-initiated transition used during planned maintenance windows. This approach provides a controlled, graceful handoff from the active node to the standby node, allowing administrators to drain running jobs, verify standby readiness, and coordinate the transition with change management processes.

A planned switchover follows this general workflow:

1. **Verify standby health**: Confirm that the standby node's Patroni, etcd, and keepalived services are running and that PostgreSQL replication is in sync with zero lag.
2. **Drain active jobs**: Disable all scheduled jobs or wait for currently running jobs to complete. This prevents job interruption during the transition.
3. **Initiate the switchover**: Use the Veeam Backup & Replication console or the `patronictl switchover` command to trigger the planned failover. Patroni will demote the current leader and promote the standby in a coordinated sequence.
4. **Validate the new active node**: Confirm that the VIP has migrated, Veeam services are accepting connections, and the configuration database is accessible on the new leader.
5. **Perform maintenance on the former active node**: With the standby now serving as the active node, the original active node can be patched, rebooted, or otherwise serviced. Once maintenance is complete, the node will rejoin the cluster as a standby.

Tip

Always perform a planned switchover before applying operating system patches or firmware updates to the active node. This practice ensures that backup operations

> continue uninterrupted during the maintenance window and that you have a known-good fallback if the maintenance procedure encounters issues.

Monitoring failover events

Effective HA management requires visibility into failover events as they occur. Veeam v13.0.1 provides multiple monitoring surfaces that administrators should configure proactively, rather than relying solely on post-incident investigation.

Veeam ONE alarms and actions

The HA status can be monitored using alarms included in Veeam ONE. The alarms you can use are:

- Veeam Backup & Replication HA cluster created
- Veeam Backup & Replication HA cluster disassembled
- Veeam Backup & Replication HA cluster primary node state
- Veeam Backup & Replication HA cluster secondary node state
- Veeam Backup & Replication HA cluster failover state
- Veeam Backup & Replication HA cluster switchover state

The alarm for the primary node state can also be used to complete automatic failover actions when the primary node becomes unavailable.

System logs and Patroni logs

For detailed troubleshooting, the Patroni and etcd logs on each cluster node provide granular visibility into the decision-making process during failover events. Key log locations on the Veeam Software Appliance include the following:

Component	Log location / command
Patroni	`journalctl -u patroni`
etcd	`journalctl -u etcd`
keepalived	`journalctl -u keepalived`
Cluster status	`patronictl -c /etc/patroni/patroni.yml list`

Technet24

Component	Log location / command
Replication lag	`patronictl -c /etc/patroni/patroni.yml list` (Lag column)

Table 7.7: Log details for Patroni and system

When investigating a failover event, start by examining the Patroni logs on the node that was promoted. Look for entries containing `acquired session lock` and `promoted self to leader` to confirm the promotion sequence. Then cross-reference with the etcd logs to verify that the leader lock was properly released and re-acquired.

Post-failover operations

After a failover event, whether automatic or planned, a series of validation and recovery tasks should be performed to ensure the cluster is fully operational and protected against a subsequent failure.

Validating active node

Immediately after failover, verify the following:

- The Virtual IP address responds to ping and is bound to the correct network interface on the new active node.
- The Veeam Backup Service is running and accepting console connections.
- The configuration database is accessible and reports the expected number of managed backup infrastructure components.
- Scheduled jobs are firing at their expected times, and any jobs that were interrupted during the failover have been restarted or rescheduled as appropriate.

Recovering a failed node

Once the new active node is confirmed operational, attention turns to recovering the failed node and restoring the cluster to its full two-node (or multi-node) configuration. The recovery process depends on the nature of the original failure:

Failure scenario	Recovery procedure
Transient failure (service crash, brief network outage)	Restart the Patroni service on the failed node. Patroni will automatically rejoin the cluster as a standby replica and begin streaming replication from the current leader.

Failure scenario	Recovery procedure
Node reboot (planned or unplanned)	After the node completes its boot sequence, Patroni and etcd services should start automatically. Verify that the node appears as a replica in the `patronictl list` output and that replication lag is converging to zero.
Corrupted database or unrecoverable failure	Reinitialize the standby node by running `patronictl reinit` to perform a fresh base backup from the current leader. This process replaces the local database entirely and is appropriate when the standby's data directory is damaged.
Hardware replacement	Deploy a fresh VSA node, configure it with the same HA cluster settings, and join it to the existing cluster. Patroni will handle the initial base backup and begin streaming replication automatically.

Table 7.8: Failure scenarios for HA

Note

WARNING!

Do not leave the HA cluster running in a degraded single-node state for extended periods. While the active node can operate independently, any subsequent failure will result in a complete service outage with no automatic recovery path. Prioritize restoring the standby node to reestablish full HA protection.

Failover testing and validation

An HA deployment that has never been tested under realistic failure conditions provides a false sense of security. Veeam recommends conducting regular failover drills to validate the cluster's behavior and to ensure that operations teams are familiar with the failover and recovery procedures.

Technet24

Recommended testing cadence

At a minimum, conduct the following tests on a quarterly basis or after any significant infrastructure change (such as network reconfiguration, firmware updates, or the addition of new managed components):

- **Planned switchover test**: Initiate a manual switchover and verify that the standby node assumes the active role cleanly, all services restart, and the VIP migrates correctly.
- **Simulated active node failure**: Shut down the active node abruptly (power off, not graceful shutdown) and observe the automatic failover behavior. Measure the actual failover time and compare it against your organization's RTO requirements.
- **Network partition test**: Temporarily disconnect the network link between the two HA nodes and verify that the cluster handles the split-brain scenario correctly, with only one node assuming the active role.
- **Job continuity validation**: Run a backup job during the failover test and confirm that the job either resumes from its checkpoint on the new active node or is correctly flagged for manual restart.
- **Failback and recovery**: After each failover test, recover the original node and verify that it successfully rejoins the cluster as a standby with zero replication lag.

Tip

Document the results of each failover test, including the measured failover time, any anomalies observed, and the steps taken during recovery. This documentation is invaluable for audit and compliance purposes, and it builds institutional knowledge that accelerates response during actual outage events.

Common failover issues and troubleshooting

While the HA failover mechanism in Veeam v13.0.1 is designed to be reliable and self-healing, certain environmental conditions can interfere with clean failover operations. The following table summarizes common issues and their resolutions:

Symptom	Likely cause	Resolution
Failover does not occur after active node failure	etcd cluster has lost quorum due to multiple node failures or misconfigured endpoints.	Verify etcd cluster health with `etcdctl` endpoint health. Restore quorum by recovering at least a majority of etcd members.
VIP does not migrate after Patroni promotion	keepalived service not running or misconfigured VRRP instance on the standby node.	Check keepalived status and configuration. Verify the VRRP `virtual_router_id` matches across all nodes and that the check script correctly detects the Patroni leader.
Standby node shows persistent replication lag	Network bandwidth constraints, high write volume on the active node, or slow disk I/O on the standby.	Monitor network throughput between nodes. Check disk I/O metrics on the standby. Consider enabling WAL compression if bandwidth is limited.
Split-brain: both nodes claim the active role	Network partition with both nodes retaining etcd connectivity to different portions of the cluster.	This should not occur in a properly configured two-node cluster with a single etcd endpoint. If it does, immediately stop Patroni on one node and investigate the etcd topology.
Veeam services fail to start after promotion	Database promotion completed, but Veeam Backup Service failed to detect the change or encountered a configuration mismatch.	Restart the Veeam Backup Service manually. Review the Veeam service logs for connection errors or configuration inconsistencies.

Table 7.9: Common issues and troubleshooting

Technet24

Best practices summary

To maximize the reliability and effectiveness of HA failover in Veeam Backup & Replication v13.0.1, observe the following best practices:

- **Deploy HA nodes on physically separate infrastructure**: Use independent servers, network switches, and power circuits to eliminate common points of failure.
- **Monitor replication lag continuously**: A standby with significant lag may result in data loss (transactions committed on the active node but not yet replicated) during an unplanned failover.
- **Use planned switchovers for all maintenance activities**: Never perform updates or patches directly on the active node without first failing over to the standby.
- **Test failover regularly**: Conduct quarterly failover drills and document the results for operational readiness and compliance purposes.
- **Restore degraded clusters immediately**: Do not operate in a single-node degraded state any longer than necessary. A degraded cluster offers no failover protection.
- **Configure alerting for HA events**: Set up email or SNMP notifications for Patroni leadership changes, replication lag thresholds, and keepalived state transitions so that the operations team is immediately aware of any cluster state changes.
- **Document your HA topology**: Maintain an up-to-date diagram showing node roles, network configuration, VIP assignment, and etcd membership. This documentation is critical for onboarding new team members and for rapid incident response.

HA failover is a foundational capability of the Veeam Software Appliance architecture in v13.0.1. When properly configured, monitored, and tested, it ensures that backup infrastructure remains available even in the face of hardware failures or maintenance requirements. In the next section, we will explore how centralized update management integrates with HA to deliver zero-downtime patching across the backup infrastructure.

Summary

This chapter has provided you with information regarding the new High Availability (HA) clustering feature of Veeam Backup & Replication v13. We discussed what HA is within Veeam Backup & Replication v13 and how to configure it. Then, we addressed how failover works for both manual and automatic scenarios, ensuring data resiliency for your HA cluster. Lastly, we examined how HA ensures you have a disaster recovery plan that will ensure your jobs continue to run even with an unexpected outage.

In the next chapter, we will cover the new enhancements for Continuous Data Protection (CDP), including the new Universal CDP for Veeam Backup & Replication v13.

Further reading

- High Availability (HA) Cluster - `https://helpcenter.veeam.com/docs/vbr/userguide/high_availability_cluster.html?ver=13`
- How High Availability (HA) Cluster works - `https://helpcenter.veeam.com/docs/vbr/userguide/high_availability_hiw.html`
- Infrastructure for High Availability Cluster - `https://helpcenter.veeam.com/docs/vbr/userguide/high_availability_infrastructure.html?ver=13`
- High Availability – Considerations and Limitations - `https://helpcenter.veeam.com/docs/vbr/userguide/high_availability_limitations.html?ver=13`

Technet24

8

Veeam CDP – Universal Continuous Data Protection

Veeam continuous data protection (CDP) is part of the **Veeam Backup & Replication** but has been enhanced and now has a **Universal CDP** option. This chapter will discuss what is new in CDP, including the new Universal CDP. We will look at how CDP can be used to improve network resilience for recovery. Lastly, we will explore the new enhanced I/O journal that is used with CDP, as well as other enhancements.

In this chapter, we're going to cover the following main topics:

- What's new with Universal CDP in v13.0.1
- Network resilience with CDP
- Enhanced I/O Journal scalability and additional CDP features

Technical requirements

To follow along in this chapter, you should have access to the following:

- You must have installed Veeam Backup & Replication with a hypervisor configured (VMware vCenter/vSphere in the examples we will be discussing)
- I/O Filter driver installed into your hypervisor cluster: `https://helpcenter.veeam.com/docs/vbr/userguide/cdp_io_filter_install.html?ver=13`
- CDP Proxies deployed: `https://helpcenter.veeam.com/docs/vbr/userguide/cdp_proxy.html?ver=13`

Technet24

What's new with Universal CDP in v13.0.1

Continuous data protection (CDP) has always been one of Veeam's most powerful capabilities for protecting Tier-1 workloads. Version 13.0.1 takes a significant leap forward by removing the hypervisor dependency entirely and introducing **Universal CDP**, an agent-based replication engine that can protect any Windows machine, regardless of whether it lives on vSphere, Hyper-V, bare metal, or in the cloud.

The vSphere limitations

Prior to version 13, CDP in Veeam Backup and Replication was exclusively tied to VMware vSphere. The technology relied on vSphere APIs for I/O Filtering (VAIO) and required a dedicated I/O filter to be installed directly on each vSphere cluster that contained protected VMs. While this delivered outstanding results for VMware environments (near-zero RPOs, no snapshots, replicas in a ready-to-start state), it left a significant gap: physical servers, cloud instances, and workloads running on non-VMware hypervisors had no equivalent protection path.

For organizations running mixed workloads, this meant that critical Windows servers on Hyper-V, physical database servers, or cloud VMs simply could not benefit from the same continuous replication that their vSphere counterparts enjoyed. Version 13.0.1 closes that gap with the introduction of Universal CDP.

What is Universal CDP?

Universal CDP is an agent-based continuous replication capability that decouples CDP from the hypervisor layer entirely. Rather than intercepting I/O at the vSphere kernel layer via VAIO, Universal CDP installs two lightweight components directly on the source workload:

- **Veeam CDP agent service**: Deploys the volume filter driver, receives the captured I/O data, aggregates it, and forward it to the CDP proxies
- **Veeam CDP volume filter driver**: A kernel-level driver that intercepts I/O transactions generated by the operating system and passes them to the Agent Service in real time

This architecture means the source workload does not need to reside on vSphere at all. It only needs to be a supported operating system with the two agent components installed. In v13.0.1, Universal CDP supports Windows workloads, covering physical servers, virtual machines running on any hypervisor, and cloud instances.

Note

New in v13.0.1:

Universal CDP for Windows machines (physical, virtual, and cloud) targeting VMware vSphere and VMware Cloud Director. Linux support and additional replication targets are on the roadmap for subsequent releases.

How it works

Ever wondered how the CDP service works? Well, let's take a look at the overall components and how you leverage CDP.

CDP coordinator service

The Veeam **CDP coordinator service** is the central orchestration engine for Universal CDP. When a Universal CDP policy starts, the coordinator performs the following sequence before data transfer begins:

1. Queries the Veeam CDP agent service on the source workload for detailed information about the guest OS, hardware configuration, network adapters, and full storage layout.
2. The agent service transforms the raw storage layout into a synthetic disk layout that can be recreated on the VMware target.
3. The coordinator requests the target vCenter Server to create replica VMs with empty disks, built from the synthetic layout.
4. The coordinator requests the agent service to attach the volume filter driver to each relevant volume on the source machine.
5. Data transfer begins immediately after the initial configuration is complete.

Once active, the coordinator continues to manage RPO settings, short-term and long-term retention policies, CDP proxy selection, and network throttling rules. It functions as the single control plane for the entire protection lifecycle.

An important resilience characteristic of this is that if the coordinator service goes offline, existing CDP policies continue to operate and maintain their short-term restore points. Long-term restore points will not be created or pruned until the coordinator is restored, but active replication is not interrupted.

Volume filter driver: I/O interception without snapshots

The volume filter driver is the key innovation that allows Universal CDP to achieve near-zero RPO without creating snapshots. The driver operates at the kernel level and intercepts every I/O write operation as it transits between the application layer and the storage subsystem. These

Technet24

intercepted writes are passed to the CDP agent service, which batches and transmits them to the source CDP proxy. The proxy then forwards them onward to the target infrastructure, where they are applied continuously to the CDP replica.

Because there are no snapshots involved, there is no periodic I/O consolidation overhead on the source workload, no snapshot growth risk, and no snapshot commit window. The RPO is constrained only by network latency and throughput between source and target. In most LAN environments, this translates to single-digit second RPOs.

Short-term and long-term restore points

Universal CDP uses the same dual restore-point model familiar from CDP for VMware vSphere:

- **Short-term restore points** are maintained continuously in a rolling journal. These allow recovery to any point within a configurable window measured in seconds or minutes, effectively delivering near-zero RPO. The journal retains short-term restore point data for a maximum of 168 hours (7 days).
- **Long-term restore points** are created at administrator-defined intervals and retained for a configurable period measured in hours or days. These provide a rollback capability for scenarios such as ransomware or logical data corruption that may not be detected immediately.

The coordinator service manages the creation and lifecycle (pruning) of both restore point types. Short-term points are created and removed autonomously by the CDP infrastructure even if the coordinator is temporarily unavailable; long-term point management requires the Coordinator to be online.

Replication targets in v13.0.1

The initial release of Universal CDP supports two replication target types: VMware vSphere and VMware cloud director.

VMware vSphere

Replicas are created and maintained as standard VMware virtual machines in a ready-to-start state on the target vSphere cluster. The target vSphere environment requires the Veeam CDP I/O filter to be installed on the cluster that will host replica VMs. The I/O filter on the target side is responsible for managing the replica disk data and maintaining the short-term restore point journal. Failover to a vSphere target uses the standard CDP failover workflow and brings the replica VM online within seconds.

VMware cloud director

Universal CDP to VMware Cloud Director is particularly significant for service providers and cloud-first organizations. In this model, source Windows workloads (physical, virtual, or

cloud) are continuously replicated into organization VDCs within a Cloud Director environment. The unit of management on the target side aligns with Cloud Director constructs (vApps and the VMs within them), rather than individual standalone VMs.

The I/O filter must be installed on the VMware vSphere cluster that backs the organization VDC used for CDP replicas. Failover in a cloud director context triggers the cloud director vApp failover workflow, bringing protected workloads online within the organization VDC.

Integration with Veeam cloud connect for DRaaS

Universal CDP integrates directly into the existing Veeam cloud connect replication infrastructure, which has substantial implications for service providers. Service providers already operating a cloud connect replication environment with cloud gateways, hardware plans, and organization VDCs already configured can extend those same resources to offer Universal CDP as a managed DRaaS service to tenants, without a separate infrastructure build-out.

CDP in the cloud connect environment maintains the same multi-tenant architecture as snapshot-based cloud connect replication. CDP infrastructure components are distributed between the tenant side and the service provider side, and all communication flows through the existing cloud gateway layer using persistent, encrypted connections. The Veeam CDP coordinator services on both the SP backup server and the tenant backup server communicate through this gateway to coordinate replication tasks and control resource allocation.

For tenants, the experience is straightforward: deploy the CDP agent service and volume filter driver on the source Windows workloads, connect to the service provider via the standard Cloud Connect registration process, and then create a Universal CDP policy targeting the cloud host provided by the SP. The resulting CDP replicas live within the SP's infrastructure and can be failed over on demand.

Deployment overview

The following steps summarize the setup sequence for Universal CDP in an on-premises deployment:

1. **Installing CDP agent components**: Deploy the Veeam CDP agent service and Veeam CDP volume filter driver on each source Windows workload to be protected.
2. **Deploying source CDP proxy**: Add a CDP proxy on the source side to handle data aggregation and transmission from the source workloads.
3. **Preparing target infrastructure**: Install the Veeam CDP I/O filter on the target VMware vSphere cluster (or cloud director organization VDC) where replicas will reside.

Technet24

4. **Deploying target CDP proxy**: Add a CDP proxy on the target side to receive and apply incoming I/O data to the replica disks.
5. **Creating a Universal CDP policy**: Define the source workloads, replication target, RPO settings, and short-term/long-term retention parameters. The coordinator performs initial synchronization and then maintains continuous replication.

Licensing

Universal CDP capabilities are available as part of the Veeam data platform and are covered under the **Veeam universal license (VUL)** model. This means existing VUL entitlements can cover Universal CDP workloads without separate or specialized licensing, simplifying capacity planning for organizations protecting mixed environments. The feature requires the VDP advanced or premium edition; it is not available under the foundation edition. Service providers should refer to the Veeam cloud and service provider program documentation for consumption and billing details specific to Universal CDP delivered as DRaaS.

Current limitations and known constraints

As of v13.0.1, the following constraints apply to Universal CDP:

- **Windows only**: The source workload agent currently supports Windows operating systems only. Linux support is on the roadmap for subsequent release.
- **Target scope**: Supported replication targets are VMware vSphere and VMware Cloud Director. Support for additional targets (such as Microsoft Azure, Nutanix AHV, and others) is planned for future releases.
- **Network requirements**: A minimum of 100 Mbps network bandwidth is recommended between source and target CDP proxies. Slower networks may be unable to sustain continuous replication without interruption under write-heavy workloads. It is recommended, however, to have 1 or 10GB/s network speeds to make the most of CDP.
- **Maximum CDP policies**: A single backup server supports a maximum of 400 CDP policies (shared with CDP for VMware vSphere).
- **Coordinator dependency for long-term points**: Long-term restore point creation and pruning requires the CDP coordinator service to be online. Short-term replication continues independently.

The following table summarizes the features of Universal CDP:

Feature	Detail
Source workload types	Physical servers, VMs (any hypervisor), cloud instances

Feature	Detail
Source OS support (v13.0.1)	Windows (Linux on roadmap)
Replication targets (v13.0.1)	VMware vSphere, VMware Cloud Director
I/O capture method	Agent-based kernel-level volume filter driver (no snapshots)
RPO capability	Near-zero (seconds to minutes, configurable)
Restore point types	Short-term (journal, up to 7 days) + Long-term (scheduled)
Cloud Connect support	Yes — uses existing VCSP Cloud Connect Replication infrastructure
Licensing	Veeam Universal License (VUL) — VDP advanced or premium
Max CDP policies per server	400 (shared with CDP for VMware vSphere)

Table 8.1: Universal CDP features

The v13.0.1 release of Universal CDP is deliberately scoped as a foundation release. By delivering the agent-based I/O interception architecture, Cloud Director target support, and cloud connect integration in this version, Veeam has established the core technical infrastructure upon which additional platform support can be layered rapidly. Linux agent support and additional replication target types are the most immediately anticipated expansions.

For enterprise architects and service providers, the significance of Universal CDP extends beyond its current feature set. It represents a fundamental shift in how Veeam approaches continuous protection: from a vSphere-centric capability to a platform-agnostic data protection primitive. As the target list expands to include public cloud platforms, additional hypervisors, and potentially even object storage, Universal CDP has the architectural potential to become the standard method for protecting any Tier-1 workload, regardless of where it runs.

Technet24

Network resilience with CDP

This section covers the end-to-end configuration of CDP policies, the architectural components that make CDP function, and critically, the strategies for ensuring network resilience in environments where CDP traffic must traverse unreliable or bandwidth-constrained links.

Configuring continuous data protection

CDP represents a shift in how Veeam Backup & Replication approaches data protection for mission-critical workloads. Unlike traditional snapshot-based backup, which operates on scheduled intervals, CDP captures every write operation as it occurs on the source VM, streaming changes to a target host in near real-time. This enables recovery point objectives (RPOs) measured in seconds rather than hours, making CDP the appropriate choice for Tier-1 workloads where even minimal data loss is unacceptable.

In Veeam Backup & Replication v13.0.1, CDP leverages VMware vSphere APIs for I/O filtering (VAIO) to intercept disk write operations at the hypervisor level. The I/O filter, deployed as a VIB on ESXi hosts, operates transparently between the guest operating system and the underlying storage, capturing every block change without requiring agents inside the VM or creating application-visible snapshots. This architecture eliminates the performance penalties traditionally associated with snapshot-based replication and avoids the snapshot stun events that can impact latency-sensitive applications.

Overview of CDP architecture and components

Before configuring a CDP policy, it is essential to understand the components involved and how data flows between them. CDP introduces several infrastructure elements that do not exist in traditional backup workflows, and a clear understanding of these elements is necessary to troubleshoot performance issues and design resilient topologies.

The following figure depicts the components needed to ensure Veeam CDP works as expected within your environment.

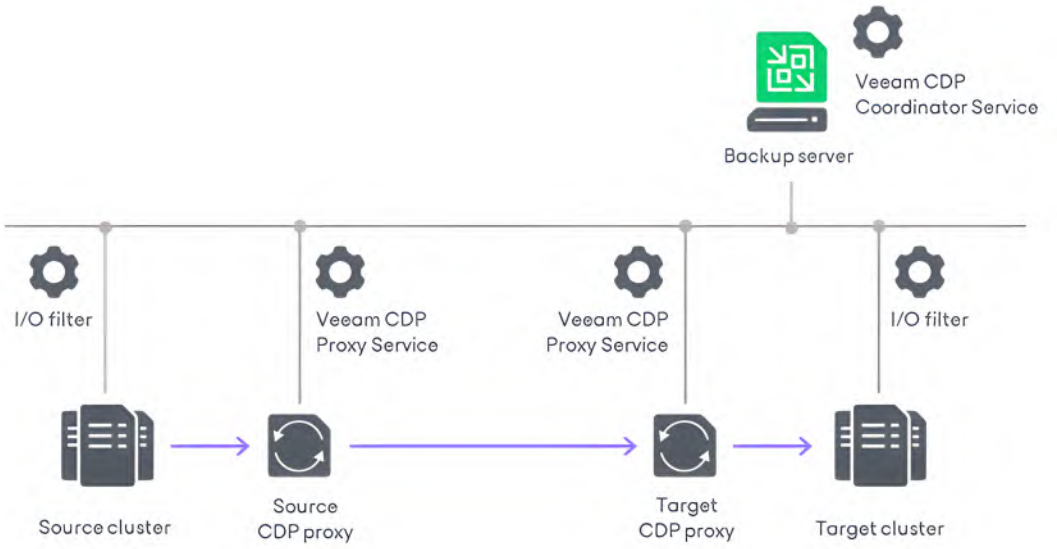


Figure 8.1: Veeam CDP architecture

Here, you can see the CDP coordinator service on the main backup server, CDP proxy servers on both the source and destination clusters which run the CDP service, and lastly the I/O filter driver that is installed to each host within the cluster for source and destination.

Core components

The CDP architecture in Veeam Backup & Replication v13.0.1 comprises the following components, each serving a distinct role in the replication pipeline:

Component	Description and role
VMware I/O filter	A VIB is installed on each ESXi host participating in CDP. It intercepts write operations at the VAIO layer and streams changed blocks to the source CDP proxy. The filter operates in the I/O path between the guest OS and the storage subsystem.
Source CDP proxy	A dedicated Veeam proxy role deployed on or near the source ESXi cluster. Receives changed blocks from the I/O filter, compresses and optionally encrypts the data, and transmits it to the target CDP proxy. Must run on a Linux-based or Windows machine (physical or virtual).

Technet24

Component	Description and role
Target CDP proxy	Receives the incoming CDP stream and writes replica data to the target data store. It also runs on Linux. The target proxy manages short-term and long-term restore points and handles the application of I/O journal data to the replica VM disks.
CDP policy	The Veeam job construct that ties together source VMs, source and target proxies, target datastores, RPO settings, and network mapping rules. A single CDP policy can protect multiple VMs, but all VMs in a policy share the same RPO target and proxy assignments.
I/O journal	A circular buffer maintained on the target side that stores granular write operations. The journal enables recovery to any point in time within the configured short-term retention window. Journal data is periodically consolidated into long-term restore points (snapshots).

Table 8.2: Veeam CDP core components

Note

CDP proxies must be Linux-based, Windows virtual, or physical machines. Veeam recommends deploying proxies as virtual machines on the same cluster or within the same vSphere site as the VMs they protect to minimize latency in the I/O capture path.

Data flow sequence

Understanding the CDP data flow is critical for sizing infrastructure and diagnosing performance bottlenecks. The following sequence describes how a single write operation propagates through the CDP pipeline:

1. The guest operating system issues a write operation to a virtual disk (VMDK)
2. The VMware I/O filter intercepts the write at the VAIO layer before it reaches the underlying storage. The write is not blocked; it proceeds to storage simultaneously.
3. The I/O filter streams the changed block data to the source CDP proxy over a TCP connection.
4. The source CDP proxy compresses the data (and encrypts it if network traffic encryption is enabled) and transmits it to the target CDP proxy.

5. The target CDP proxy writes the received blocks to the I/O journal on the target datastore.
6. Periodically, journal data is consolidated into long-term restore points, creating VMware snapshots of the replica VM at defined intervals.

This pipeline operates continuously. Unlike periodic backup jobs that run and complete, a CDP policy, once started, maintains a persistent connection between source and target proxies. This persistent nature has significant implications for network design, which are addressed in detail later in this section.

Prerequisites and pre-configuration requirements

CDP has specific infrastructure requirements that must be satisfied before policy creation. Failure to meet these prerequisites will result in deployment errors or degraded performance that may not be immediately apparent.

The following are the VMware vSphere requirements:

- vSphere 7.0 or later (vSphere 7.0 U3 or later recommended for optimal VAIO filter performance).
- VMware Tools installed and running on all VMs to be protected by CDP.
- ESXi hosts must have network connectivity to the source CDP proxy on TCP ports 33032, 33033, 6182, 902, 443 (configurable).
- Hardware version 7 or later for all protected VMs (hardware version 13+ recommended.) The hardware version you are protecting must be supported on the target cluster.
- Fault tolerance (FT) must not be enabled on protected VMs, as VAIO filters are incompatible with FT.

I/O filter deployment

The Veeam I/O filter must be installed on every ESXi host in any cluster containing VMs that will be protected by CDP. The filter is deployed and managed through the **Veeam Backup & Replication** console and is installed as a VIB package on the ESXi host.

To install the I/O filter, navigate to **Backup Infrastructure** | **Managed Servers**, right-click the vCenter Server entry, and select **Manage I/O Filter**. Veeam will deploy the filter to all hosts in the selected clusters. A host reboot is ***not*** required for initial installation, but the filter installation does require that no CDP-protected VMs are running on the host during deployment. Veeam will place the host into maintenance mode automatically if the **Automatically place host into maintenance mode** option is selected.

Technet24

Figure 8.2 shows the **Backup Infrastructure** | **Managed Servers** section where you select the vCenter Server entry where you either right-click to see the **Manage I/O filters** or you use the toolbar button.

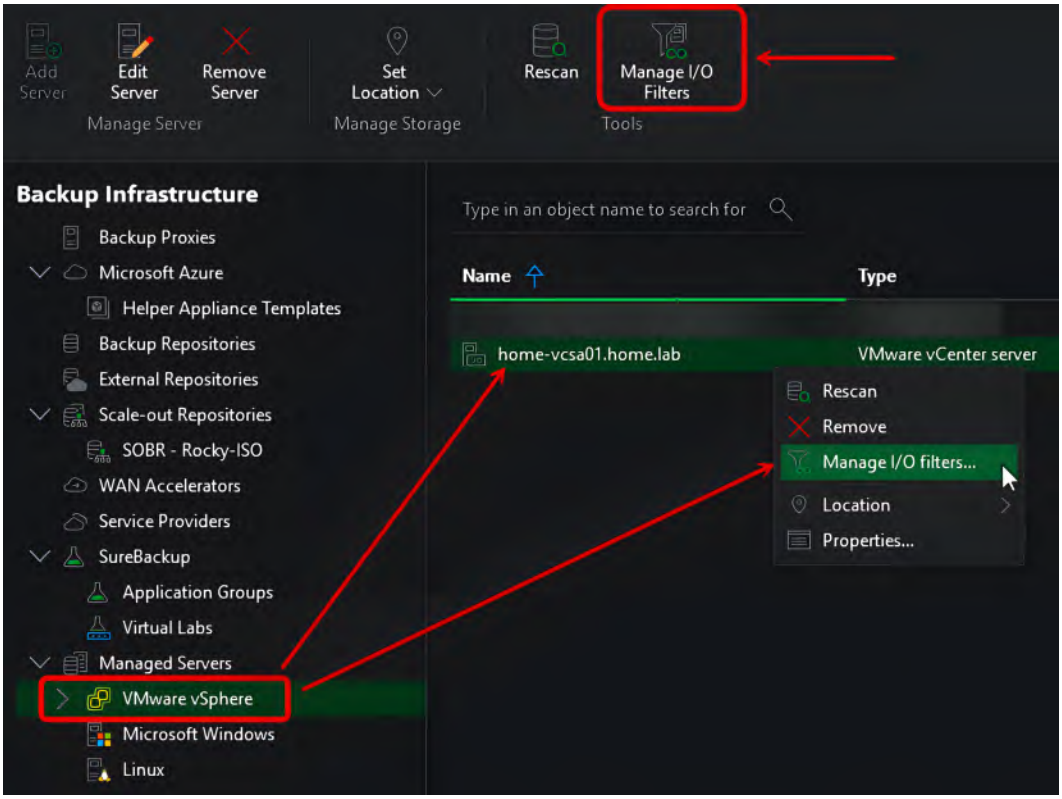


Figure 8.2: Manage I/O filters selection with right-click or toolbar button

This figure shows how to access the **Manage I/O** filters option, and *Figure 8.3* shows the deployment wizard to deploy the I/O filter to the cluster, which you click through to install the driver.

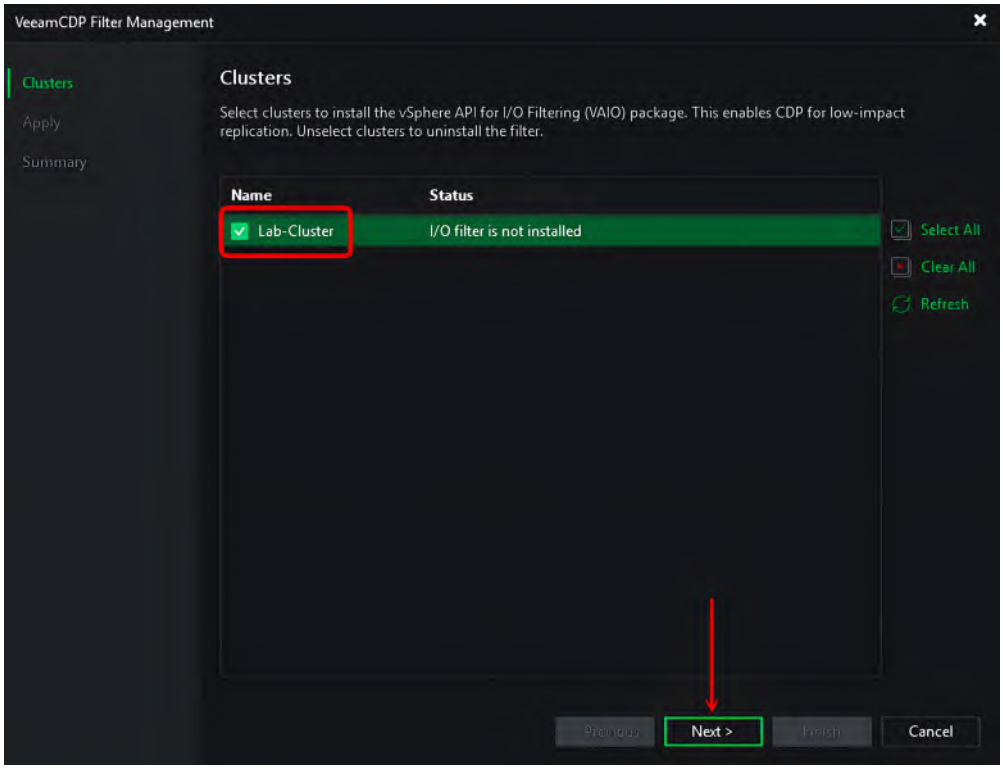


Figure 8.3: VeeamCDP filter management dialog to install I/O filters

After clicking **Next** in the above figure, you will be prompted with a question about installing the I/O filter on the selected cluster. Click **Yes** to proceed with the installation.

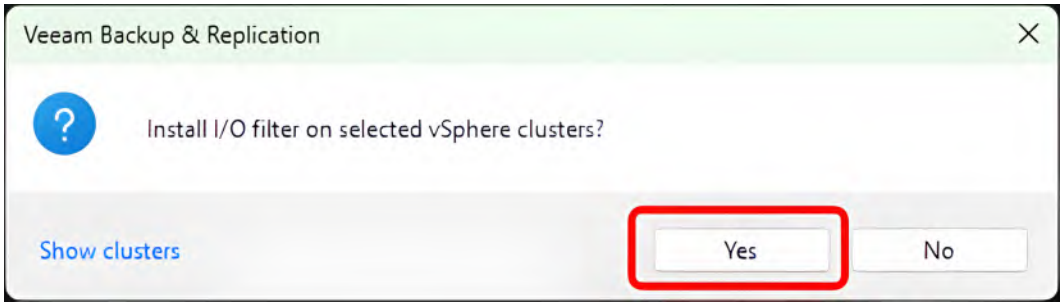


Figure 8.4: Confirming I/O filter install on selected cluster

After you click **Yes**, the installation of the I/O filter on each host commences. Once complete, you can see the details in the **Apply** window of the wizard as shown in *Figure 8.5*.

Technet24

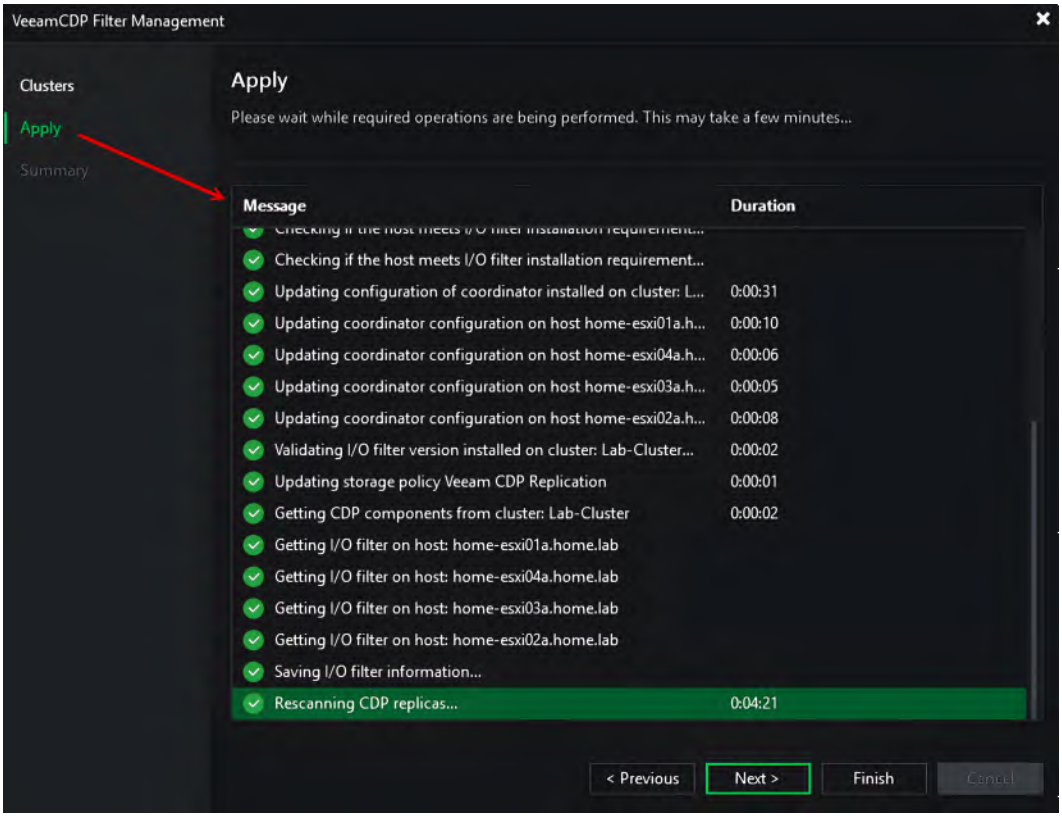


Figure 8.5: Install steps and completion for the I/O filter

Once the installation is completed, click **Next** to proceed to the **Summary** screen as shown in the following figure and click **Finish**.

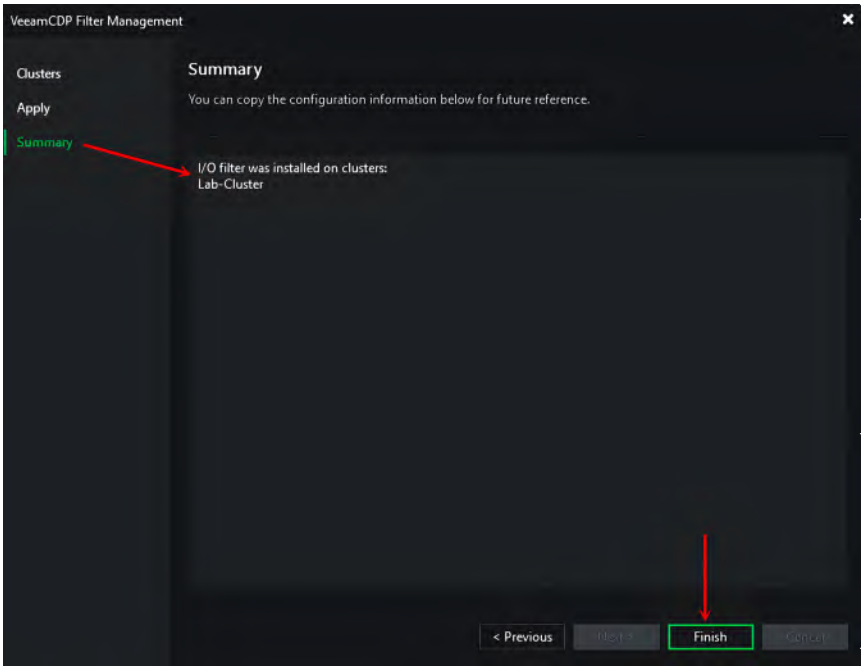


Figure 8.6: Summary of I/O filter driver installation

Note

Important:

If you are using VMware vSAN, the I/O filter installation process includes an additional step where the filter is registered as a vSAN fault domain component. Ensure your vSAN cluster has sufficient fault tolerance capacity before initiating the I/O filter deployment, as host maintenance mode transitions will temporarily reduce available capacity.

That completes the I/O filter driver installation, and we will now take a look at the CDP Proxy sizing and deployment steps in the next section.

Technet24

CDP proxy sizing considerations

CDP proxy sizing directly impacts the achievable RPO. An undersized proxy will introduce latency in the replication pipeline, causing the actual RPO to drift from the target RPO configured in the policy. The following table provides general sizing guidance based on workload characteristics:

Workload profile	Write rate	vCPU	RAM	Network
Light (file servers, web)	< 20 MB/s	4 vCPU	8 GB	1 Gbps
Medium (app servers)	20–80 MB/s	8 vCPU	16 GB	10 Gbps
Heavy (databases, OLTP)	80–200 MB/s	16 vCPU	32 GB	10+ Gbps
Extreme (large SQL, SAP)	> 200 MB/s	16+ vCPU	64 GB	25 Gbps

Table 8.3: Veeam CDP proxy sizing requirements

Using these above requirements, you will need to deploy at least two CDP proxies, one in the main site (PROD) and one in your DR site (DR). These can be deployed using the Veeam Infrastructure Appliance (VIA) ISO file using the required specifications noted above and from *Chapter 3, Just Enough Operating System (JeOS)*. Once deployed, you add the servers to your Veeam Backup & Replication 13 environment and then install the CDP proxy role as follows:

1. Click into the **Backup Infrastructure** tab and click on **Backup Proxies**.
2. In here, you can right-click or select the **Add proxy** button to create a new CDP proxy.

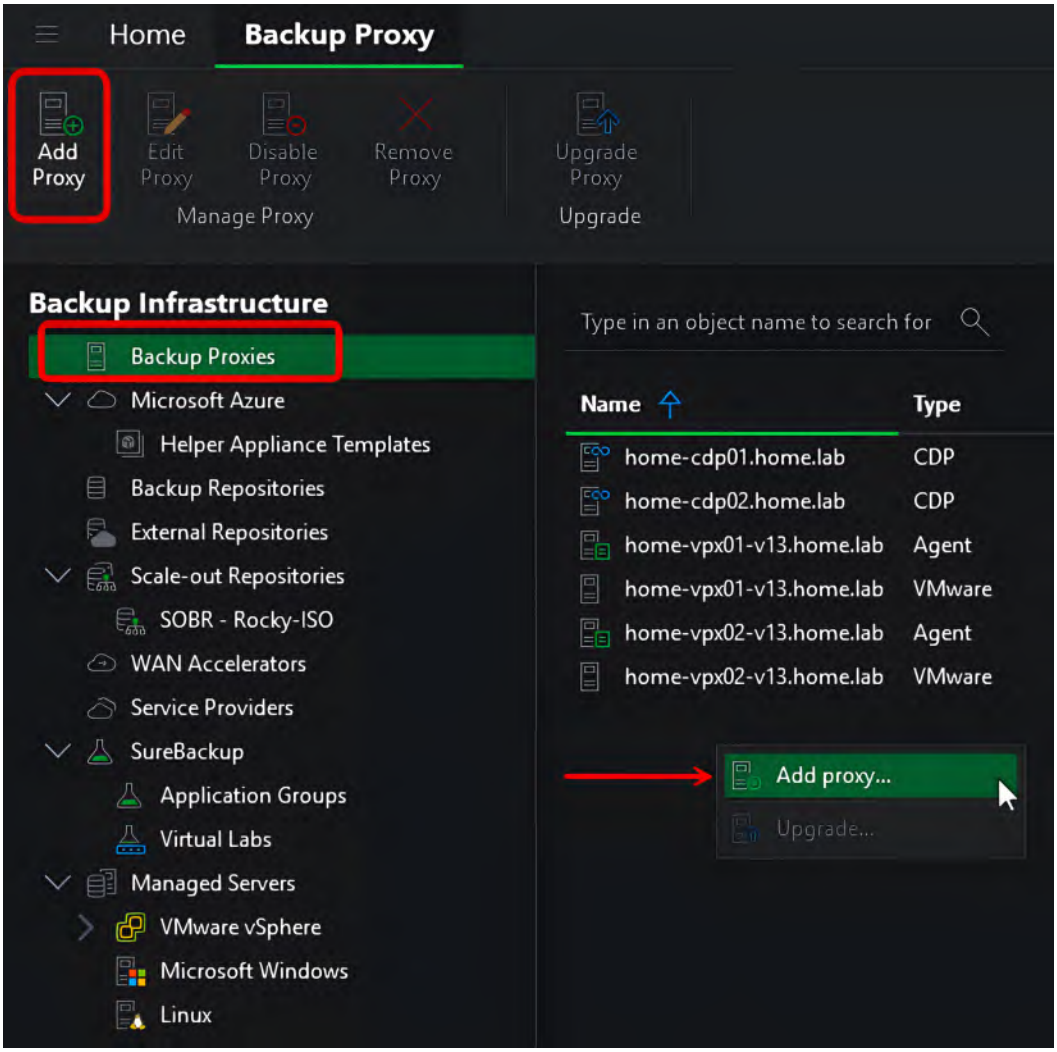


Figure 8.7: Adding a CDP proxy using right-click or toolbar button

3. In the **Add Backup Proxy** wizard, select **CDP proxy** as shown in *Figure 8.8*.

Technet24

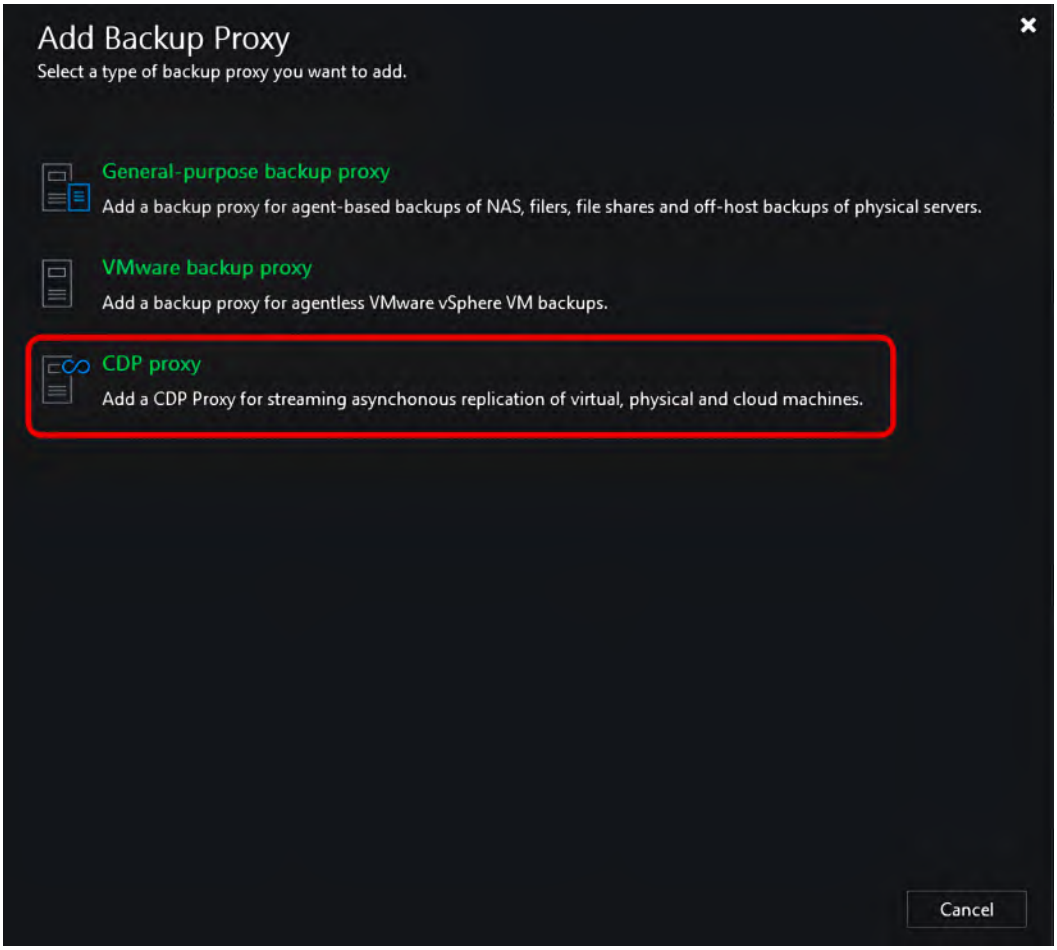


Figure 8.8: Add backup proxy – CDP proxy

4. You will now be at the **New CDP Proxy** wizard, which you will go through each step to add your CDP proxy servers (see the link in the *Further reading* section at the end of the chapter for examples of each screen)
5. Select the servers from the drop-down menu or click the **Add New...** option. Your server should already be listed after adding it to the **Managed Servers** section.

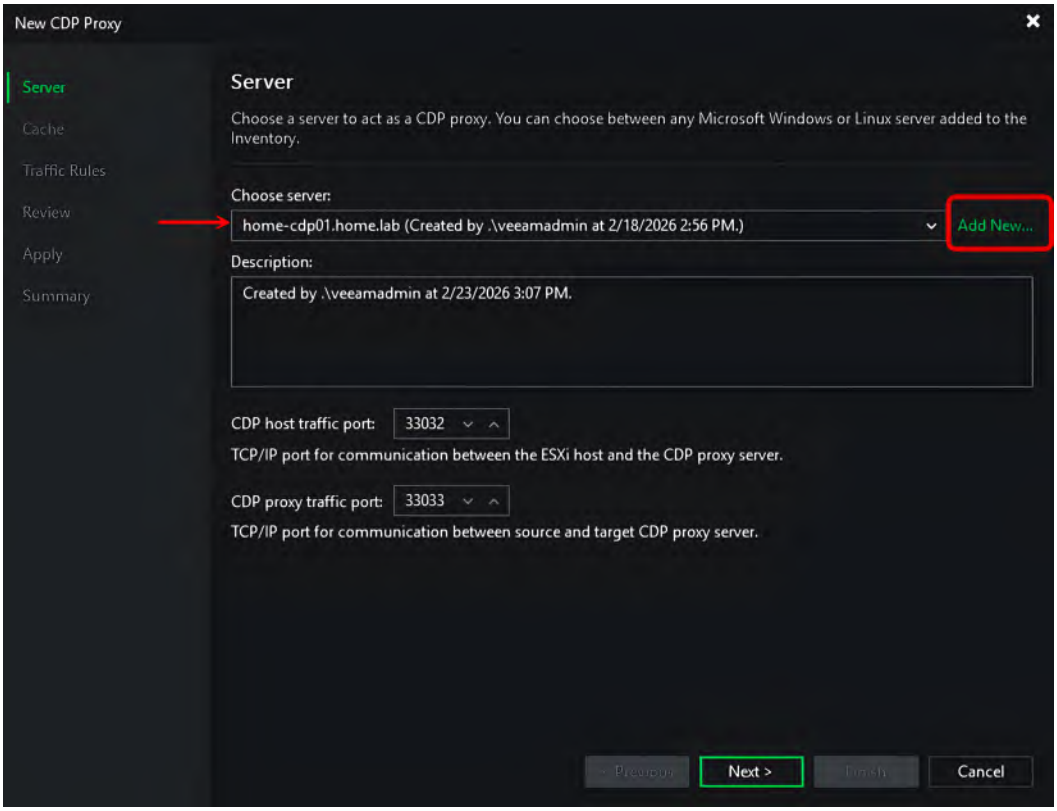


Figure 8.9: New CDP proxy wizard

This completes the deployment of your CDP proxies, which will be used in the next section for deploying CDP policies.

Tip

Use Veeam ONE or vRealize operations to baseline the write I/O rate of candidate VMs before configuring CDP. A common mistake is estimating write rates from storage performance metrics, which include read operations. CDP only replicates writes, so the write-specific throughput is the relevant metric for proxy and network sizing.

Technet24

Creating a CDP policy

With the prerequisites satisfied, I/O filters deployed, and CDP proxies provisioned, you can create a CDP policy to begin protecting workloads. The CDP policy wizard guides you through source selection, target configuration, RPO settings, and network mapping.

Step 1: Launching the CDP policy wizard

Follow these steps to launch the CDP policy wizard:

1. In the **Veeam Backup & Replication** console, navigate to **Home** | **CDP Policy** and click either **VMware vSphere** or **Universal CDP** from the ribbon.
2. Provide a descriptive name and an optional description.
   Use a naming convention that identifies the workload tier, source site, and target site (for example, `CDP-Tier1-SQL-SiteA-to-SiteB`). In my case, I am using VMware vSphere CDP to protect one of my domain controllers.

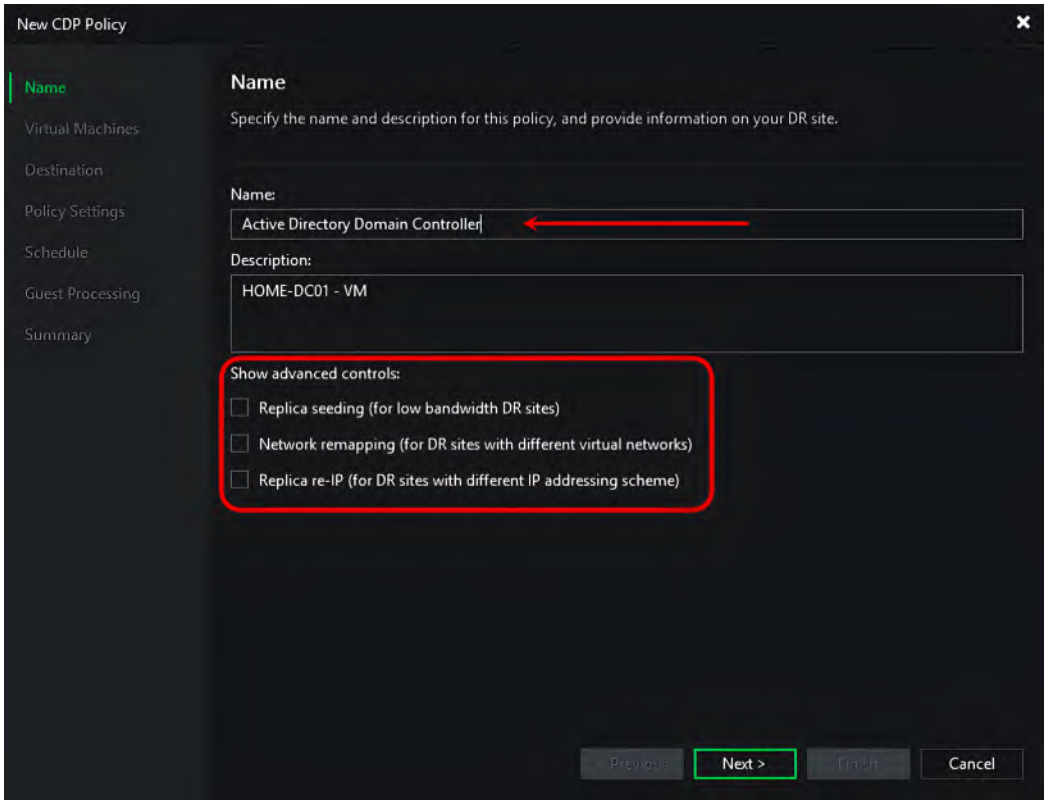


Figure 8.10: New CDP policy wizard for VMware vSphere VMs

> **Note**
>
> Universal CDP works with protection groups for physical agent servers or machines. More details can be found in the *Further reading* section at the end of the chapter.

Step 2: Selecting source VMs

1. Click **Add...** to select the VMs or VM containers (folders, resource pools, tags, or vApps) to protect.
   When adding containers, all VMs within the container are automatically included. New VMs added to the container after policy creation will be picked up on the next policy evaluation cycle. Individual VMs can be excluded from container-based selections using the **Exclusions** tab.

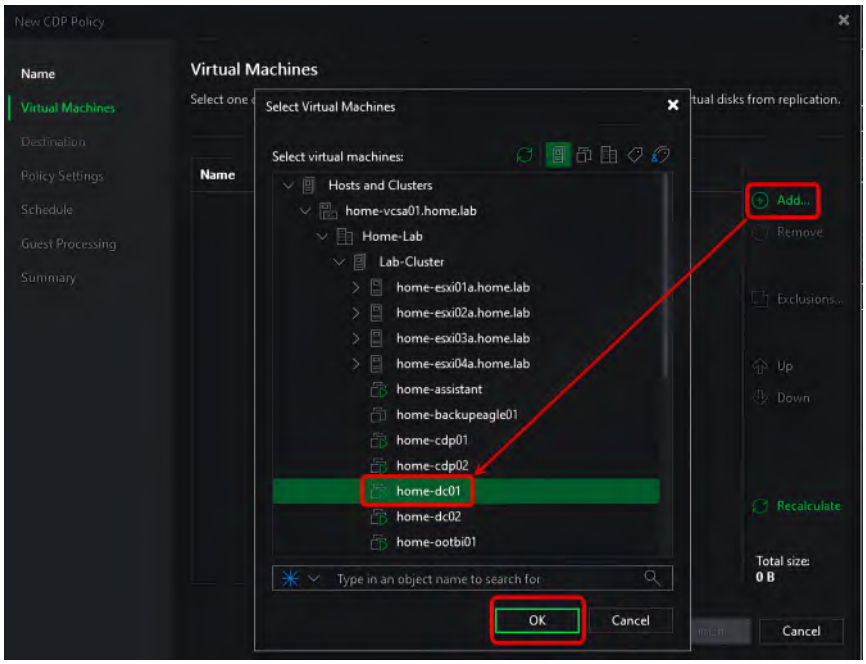


Figure 8.11: Selecting the VM(s) to protect with CDP

2. Once you select your VM(s), click **OK** as shown in *Figure 8.11* and then click **Next** on the wizard to proceed.

Technet24

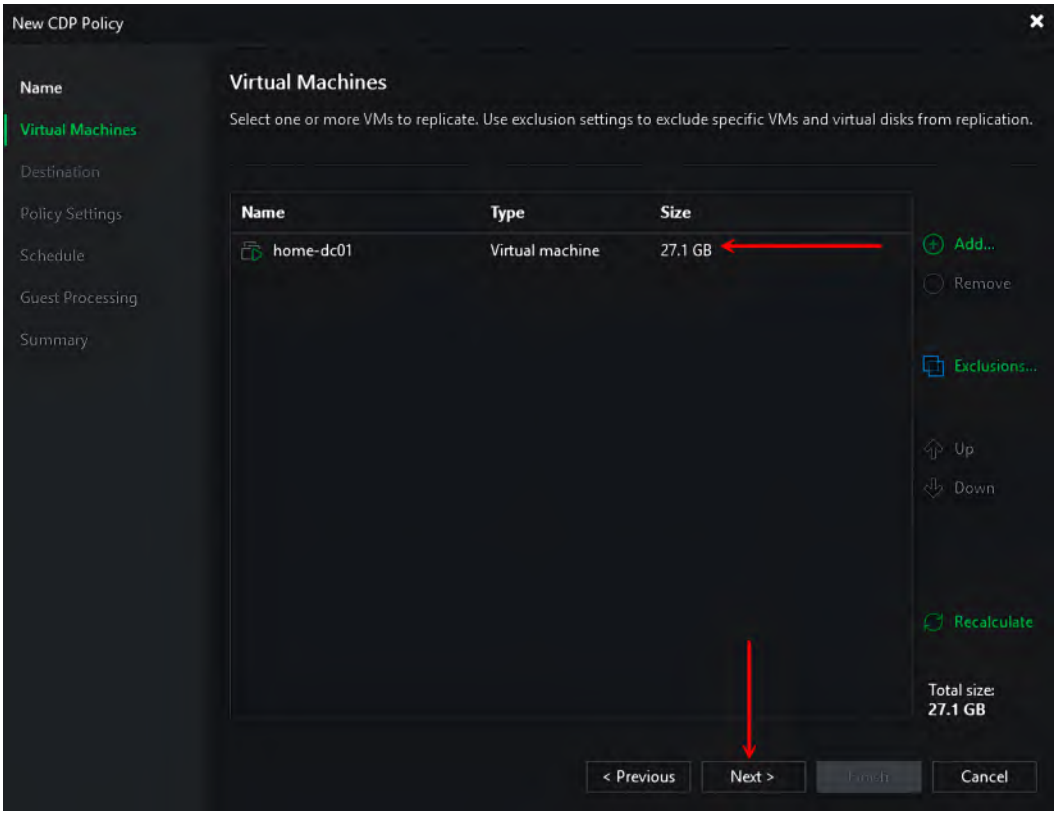


Figure 8.12: Virtual machine selection and proceeding to the next step

Step 3: Configuring the target

These steps will guide you through the process of configuring the target:

1. Select the target vCenter server, ESXi host or cluster, and datastore where replica VMs will be created. The target location should ideally be in a different failure domain from the source.
2. Configure resource pool placement, folder mapping, and datastore. You can also configure re-IP rules for failover scenarios where the target site uses a different network address space (if selected in *Figure 8.10*).

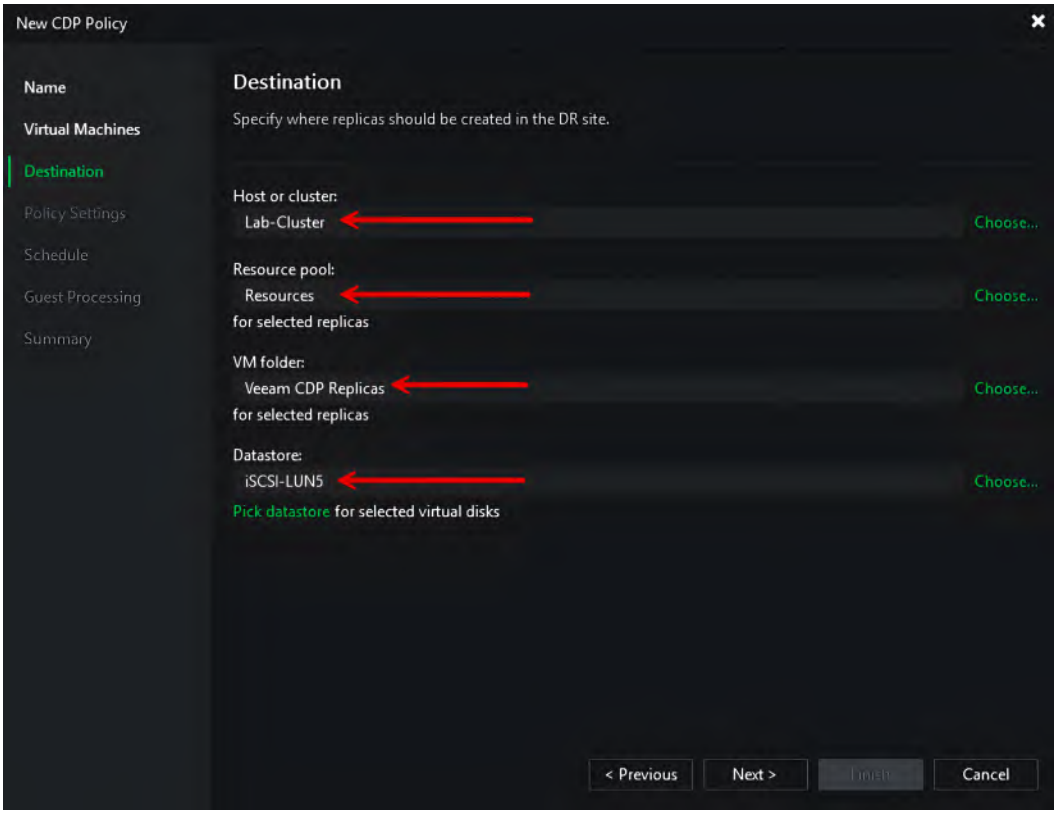


Figure 8.13: Destination for CDR replicas

Step 4: Assigning CDP proxies

Follow these steps to assign CDP proxies:

To assign CDP proxies, select the source and target CDP proxy servers. If you have deployed multiple proxies, you can assign specific proxies to the policy or allow Veeam to automatically select proxies based on proximity and load.

For production deployments, Veeam recommends explicit proxy assignments to maintain predictable performance. The VM name suffix (the default suffix is `_replica`) can be changed at this step, and I've used `_cdp_replica`.

Technet24

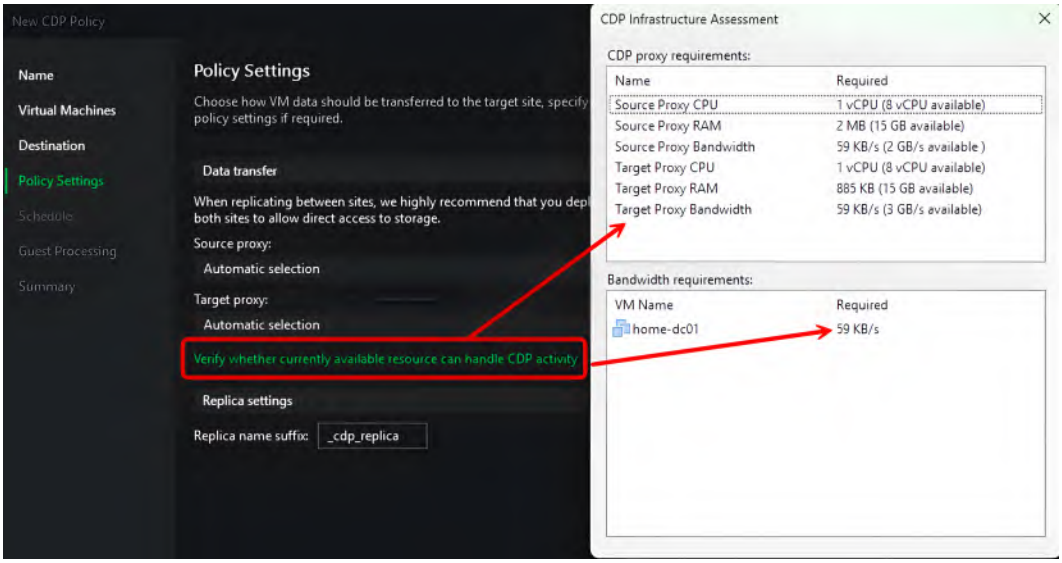


Figure 8.14: Policy settings including CDP proxy selection

Tip

If you want to verify your CDP infrastructure after you select your CDP source and target proxies, click the **"Verify whether currently available resource can handle CDP activity"**, which will then show you the requirements to complete the CDP operations, including bandwidth requirements.

Step 5: Defining RPO settings

The following steps illustrate how to define RPO settings:

1. Configure the target RPO, which determines how frequently the source CDP proxy transmits accumulated changes to the target. The minimum configurable RPO is 2 seconds, and the maximum is 24 hours.

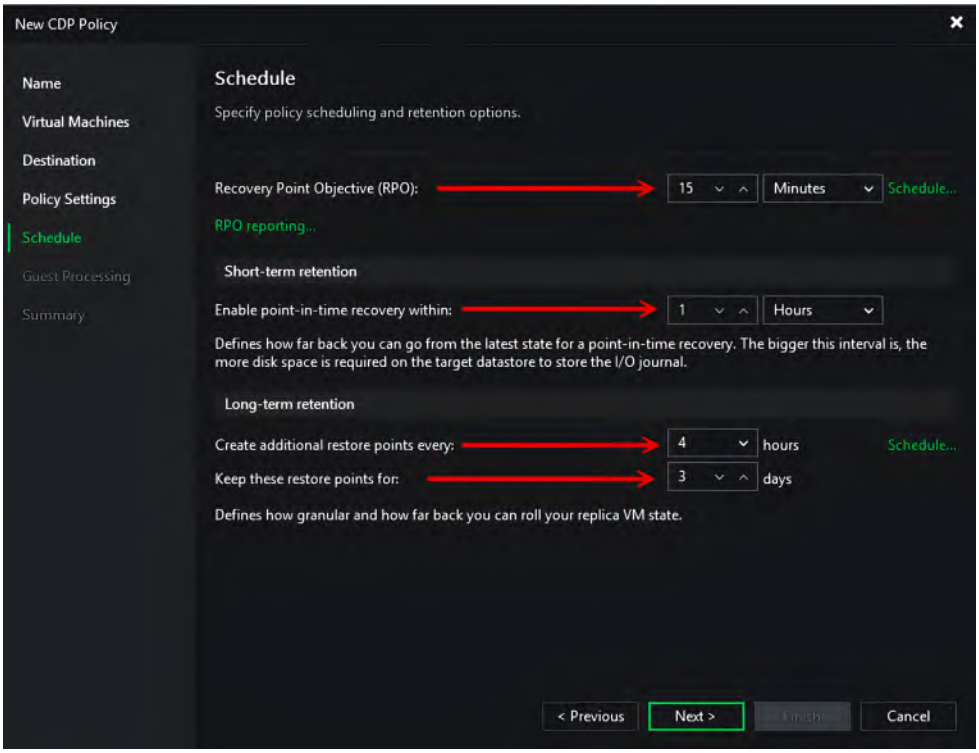


Figure 8.15: Schedule settings for RPO and short-term retention

2. Configure short-term retention (the I/O journal retention period, measured in hours) and long-term retention (the number of restore points to maintain as VMware snapshots)

RPO setting	Short-term (I/O Journal)	Long-term (Snapshots)
Granularity	Point-in-time recovery to any write operation within the journal window	Recovery to specific snapshot points only
Retention	2 to 24 hours (configurable)	1 to 28 restore points (configurable)

Technet24

RPO setting	Short-term (I/O Journal)	Long-term (Snapshots)
Storage impact	Journal size is proportional to write rate and retention window. High-write workloads with long retention will consume significant datastore space.	Each restore point consumes space equal to the changed blocks since the previous point. Impact depends on change rate.
Use case	Granular recovery from corruption, ransomware, or application errors within a recent time window.	Recovery to known-good states over days or weeks. Useful for compliance or staged recovery testing.

Table 8.4: Retention settings for continuous CDP

Step 6: Configuring network mapping

If the target site uses different network segments, configure network mapping to ensure the replica VM connects to the appropriate port groups on failover. This step is optional if source and target sites share the same network infrastructure. Network re-IP rules can also be configured for automated IP address translation during failover. You must also select this in *Figure 8.10* when you start the CDP Policy wizard to have it shown in the configuration.

Step 7: Configuring guest processing

You have the option to enable application-aware processing for your VM. Depending on the roles installed in the VM, it is recommended to enable this feature. This step is optional of course, but I am enabling it for my domain controller to ensure a successful failover in the event of a disaster.

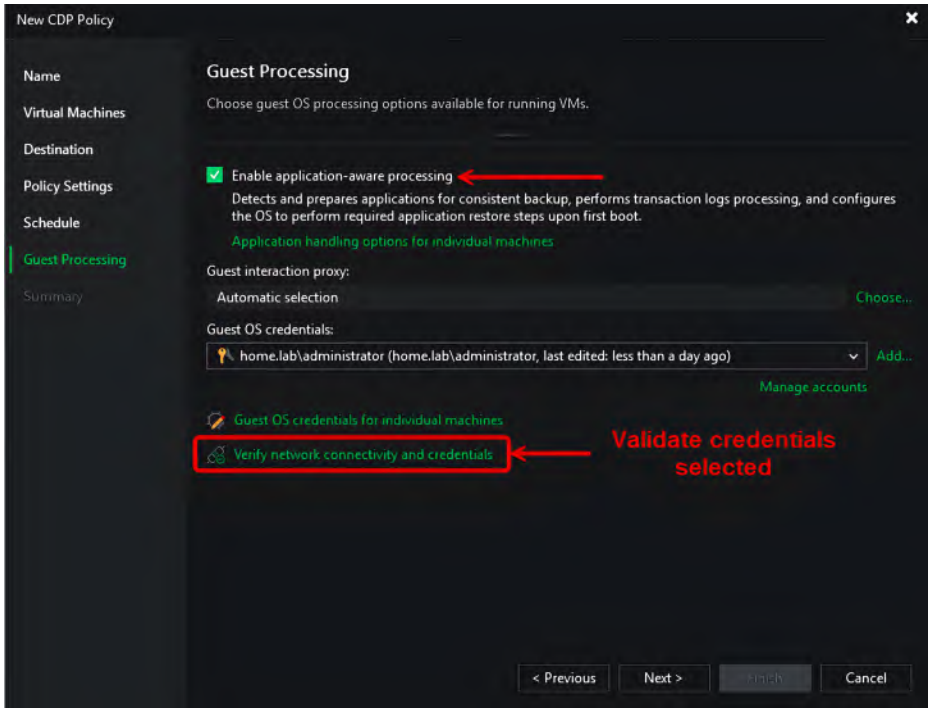


Figure 8.16: Guest processing settings – application-aware processing

Step 8: Summary and application

Review the policy summary and click **Finish** to create the policy. The policy will immediately begin an initial synchronization of the protected VMs to the target location using a full copy of the virtual disks. Once the initial sync is completed, the policy transitions to steady-state CDP operation, where only write operations are replicated.

Technet24

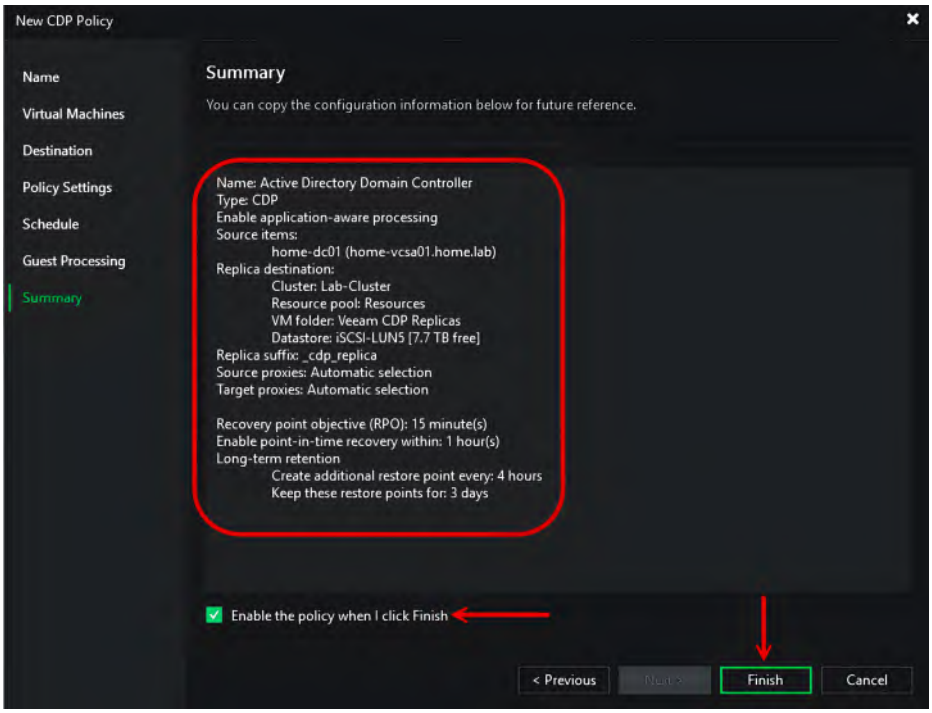


Figure 8.17: Summary of new CDP policy and enablement

Note

The initial synchronization phase can generate significant network traffic depending on the total size of the protected VMs. Plan the initial sync during maintenance windows or configure network throttling rules (discussed in the network resilience section) to prevent impact on production traffic.

The initial sync process will now run for your CDP policy job, and you can double-click the job to monitor the progress as shown in the following figure.

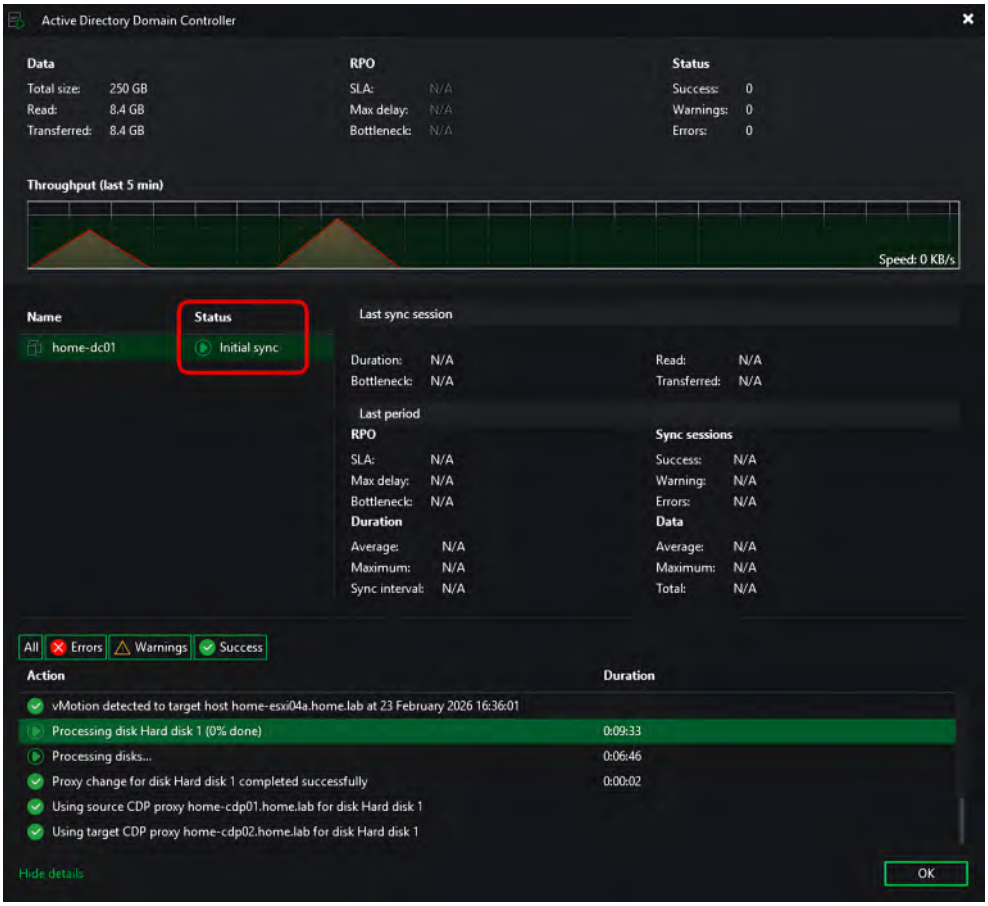


Figure 8.18: Initial sync status for CDP policy

Now that we have covered the setup and configuration of the CDP policy components and job, let's take a look at how you can improve network resilience.

Improving network resilience for CDP

The persistent, real-time nature of CDP replication places unique demands on the network infrastructure connecting source and target sites. Unlike traditional backup or replication jobs that tolerate network interruptions by simply retrying on the next schedule, CDP requires continuous connectivity to maintain its RPO guarantees. A network disruption during CDP operation does not merely delay a backup; it causes the RPO to degrade for every second the connection is unavailable, and if the disruption lasts long enough, the accumulated I/O journal on the source proxy may overflow, forcing a costly resynchronization.

Technet24

Designing for network resilience in CDP environments therefore requires a different approach than what is used for traditional backup infrastructure. This section covers the key strategies and configuration options available in Veeam Backup & Replication v13.0.1 for building resilient CDP network topologies.

Understanding CDP network failure modes

To design effective resilience, it is important to understand how CDP behaves when network connectivity is disrupted. CDP has three distinct failure response phases:

Phase	Duration	Behavior
Buffering	Seconds to minutes	The source CDP proxy buffers changed blocks in its local cache. Writes continue to be captured by the I/O filter and queued. The RPO clock is running, but data is being preserved for eventual delivery.
RPO violation	Minutes to hours	Once the buffered data exceeds the configured RPO window, the policy enters an RPO violation state. Veeam generates alerts and events. Data continues to be buffered if cache space permits. When connectivity is restored, the proxy transmits the accumulated buffer to bring the replica current.
Resynchronization	Hours	If the source proxy cache is exhausted or the disconnection exceeds the maximum cache retention, the proxy can no longer guarantee data integrity. When connectivity is restored, a full or incremental resynchronization is triggered, comparing source and replica disks block-by-block. This is resource-intensive and the policy cannot meet its RPO target until resync completes.

Table 8.5: CDP network failure modes

Note

Important:

Resynchronization is the most disruptive CDP failure mode. It requires reading and comparing all blocks on both the source and target disks, consuming significant I/O,

CPU, and network bandwidth. Design your network resilience strategy with the primary goal of avoiding resynchronization events.

Network traffic encryption and compression

CDP traffic between source and target proxies can be encrypted using AES-256 encryption. While encryption adds CPU overhead and a modest increase in latency (typically 5–15% throughput reduction depending on hardware), it is strongly recommended for any CDP traffic traversing networks outside of a fully trusted, physically secured environment. Encryption is configured at the CDP policy level under **Data Transfer** | **Encrypt network traffic**.

Compression is enabled by default and should remain enabled for most deployments. Veeam uses inline compression on the source CDP proxy before transmitting data, which reduces bandwidth consumption at the cost of source proxy CPU utilization. For workloads with highly compressible data (such as database redo logs or text-heavy application writes), compression ratios of 2:1 to 5:1 are common. For workloads that are already compressed or encrypted at the application level (such as TDE-encrypted databases), compression provides minimal benefit and the CPU cost may not be justified.

Configuring network throttling rules

In environments where CDP traffic shares bandwidth with production network flows, network throttling rules prevent CDP replication from consuming excessive bandwidth and impacting business operations. Veeam provides two mechanisms for controlling CDP network utilization.

Global Network Traffic Rules. Configured under **Main Menu** | **Network Traffic Rules**, these rules define bandwidth limits between specific network subnets. To create a rule, click **Add** and specify the source network, target network, and maximum bandwidth limit in Mbps. You can also define a schedule to apply different limits during business hours versus off-hours.

For example, to limit CDP traffic to 500 Mbps during business hours and allow unrestricted bandwidth overnight, configure as shown in the table:

Source network	Target network	Schedule	Limit
10.10.0.0/16	10.20.0.0/16	Mon–Fri 08:00–18:00	500 Mbps
10.10.0.0/16	10.20.0.0/16	All other times	Unlimited

Table 8.6: Network throttling rule examples

Technet24

Tip

When setting bandwidth limits for CDP, always account for the initial synchronization phase. The initial sync of a 2 TB VM at 500 Mbps will take approximately 9 hours. Consider temporarily increasing the limit or scheduling the initial sync during weekends to avoid prolonged RPO violations during the sync window.

Preferred networks

Veeam allows you to designate preferred networks for CDP traffic, ensuring that replication data flows over a dedicated replication VLAN or network segment rather than the general management or production network. Configure preferred networks under **Main Menu | Network Traffic Rules | Preferred Networks**. When a preferred network is configured, Veeam will route CDP traffic through interfaces attached to the preferred subnet, provided connectivity exists between source and target proxies on that network.

WAN accelerator integration

For CDP deployments that span WAN links, Veeam WAN accelerators can significantly reduce bandwidth consumption through global data deduplication. WAN accelerators maintain a persistent cache of previously transmitted data blocks, and only novel blocks are transmitted over the WAN link. This is particularly effective for workloads with repetitive write patterns, such as databases where similar structures are written repeatedly.

To configure WAN Acceleration for CDP, deploy a WAN Accelerator at both the source and target sites and assign them in the CDP policy under **Data Transfer | WAN Acceleration**. Size the WAN accelerator cache based on the total working set size of the protected VMs; Veeam recommends cache storage equal to the sum of all source VM disks multiplied by 1.1 for optimal deduplication efficiency.

Note

WAN acceleration adds latency to each data transfer operation due to the deduplication lookup process. For environments where the target RPO is under 15 seconds, test WAN acceleration carefully in a staging environment to ensure the added latency does not prevent the policy from meeting its RPO target. In very low RPO configurations (2–5 seconds), WAN acceleration is generally not recommended.

Source proxy cache configuration

The source CDP proxy maintains a local cache that buffers write data when the network connection to the target is degraded or unavailable. Proper cache sizing is the single most important factor in determining how long a CDP policy can tolerate network disruption without requiring resynchronization.

The cache size should be calculated based on the sustained write rate of the protected workloads and the maximum tolerable network outage duration. The formula is as follows:

```
Cache Size (GB)  =  Write Rate (MB/s) ×  Max Outage (seconds)  ÷  1024
```

Table 8.7 provides example calculations for common scenarios:

Write rate	Outage duration	Cache required	Recommendation
25 MB/s	30 minutes	~44 GB	Allocate 60 GB
50 MB/s	1 hour	~176 GB	Allocate 200 GB
100 MB/s	2 hours	~703 GB	Allocate 750 GB
200 MB/s	4 hours	~2.8 TB	Allocate 3 TB

Table 8.7: Example proxy cache configurations

The cache directory is configured during CDP proxy deployment. By default, Veeam places the cache on the `/var/veeam/cdp_cache` path. For production environments, this path should point to high-performance storage (NVMe SSD recommended) that can sustain the write throughput of all protected VMs simultaneously. A spinning disk cache will introduce I/O bottlenecks during both normal operation and cache drain-back after a network disruption.

Redundant network paths and proxy failover

For maximum resilience, deploy CDP infrastructure with redundant network paths between source and target sites. This can be achieved through a combination of the following strategies:

- **NIC teaming on CDP proxies**: Configure NIC teaming (bonding) on the Linux CDP proxy to aggregate multiple physical network interfaces. Use active-backup bonding mode (mode 1) for failover redundancy, or 802.3ad LACP (mode 4) for both redundancy and increased throughput. Ensure the upstream switch configuration matches the bonding mode.

Technet24

- **Multiple CDP proxies per site**: Deploy multiple source and target CDP proxies and distribute VM protection across them. If a proxy becomes unavailable, only the VMs assigned to that proxy are affected. Veeam does not currently support automatic proxy failover for CDP, so manual reassignment is required. Having pre-configured standby proxies reduces recovery time.
- **Diverse WAN paths**: For inter-site CDP, use diverse WAN circuits from different service providers. Configure routing so that CDP traffic can automatically fail over to the secondary circuit if the primary becomes unavailable. SD-WAN solutions can automate this failover and provide application-aware routing for CDP traffic.
- **Dedicated replication VLANs**: Isolate CDP traffic onto dedicated VLANs with QoS policies that guarantee minimum bandwidth. This prevents other network traffic from starving CDP replication, which is particularly important when CDP shares infrastructure with backup traffic, vMotion, or other bulk data transfer operations.

Now that we have looked at the way to enhance your network resilience, let's look into the ways you can monitor your CDP health and RPO compliance.

Monitoring CDP health and RPO compliance

Proactive monitoring is essential for maintaining CDP service levels. Veeam provides several monitoring touchpoints for CDP policies:

- **Veeam Backup & Replication console**: The CDP policy dashboard displays the current RPO status for each protected VM, including the actual versus target RPO, journal utilization, and proxy throughput. VMs in RPO violation are highlighted with warning indicators. The **Statistics** tab provides historical RPO compliance data and proxy performance metrics.
- **Veeam ONE integration**: Veeam ONE provides advanced monitoring and alerting for CDP policies, including trend analysis, capacity forecasting for journal and cache storage, and automated alerting based on configurable thresholds. Configure alerts for RPO violations, proxy disconnect events, and cache utilization exceeding defined thresholds (recommended: alert at 70% cache utilization, critical at 90%).
- **SNMP and syslog**: CDP events can be forwarded to external monitoring platforms via SNMP traps or syslog integration. This enables integration with enterprise monitoring tools such as Splunk, Grafana, or Datadog for centralized visibility across the entire infrastructure stack.

CDP network resilience best practices

The following table consolidates the key best practices discussed in this section for quick reference during deployment planning:

Practice	Recommendation
Cache sizing	Size source proxy cache to sustain the maximum tolerable network outage. Use NVMe SSD storage for the cache path. Over-provision by 20–30% above calculated requirements.
Network bandwidth	Provision 2x the sustained write rate as available bandwidth between proxies. Account for compression ratios and encryption overhead. Use dedicated VLANs with QoS guarantees.
Proxy redundancy	Deploy multiple CDP proxies per site. Distribute VMs across proxies to limit blast radius. Maintain pre-configured standby proxies for manual failover.
WAN design	Use diverse WAN circuits from different providers. Implement SD-WAN or dynamic routing for automatic path failover. Use WAN accelerators for RPO targets above 15 seconds.
Monitoring	Implement proactive alerting at 70% and 90% cache utilization. Monitor actual RPO vs. target RPO trending. Integrate with enterprise monitoring via SNMP/syslog.
Initial sync	Schedule initial synchronization during maintenance windows. Apply temporary bandwidth increases or reduce throttling during sync. Verify RPO compliance after sync completion.
Encryption	Enable AES-256 encryption for any CDP traffic traversing untrusted networks. Accept the 5–15% throughput overhead as an acceptable security trade-off. Test RPO achievability with encryption enabled.

Table 8.8: Network resilience best practices summary

By following these guidelines and properly sizing the infrastructure components, CDP can deliver near-zero RPO protection for critical workloads while maintaining resilience against the network disruptions that are inevitable in any distributed infrastructure environment. Now, let us discuss the new I/O Journal scalability and other enhancements.

Technet24

Enhanced I/O journal scalability and additional CDP features

Veeam Backup & Replication v13.0.1 delivers significant improvements to the CDP engine, building upon the scalability gains introduced in earlier releases while extending CDP capabilities to entirely new workload types. This section examines the enhanced I/O journal scalability, the evolution of the CDP I/O filter architecture, and the transformative introduction of Universal CDP as an agent-based replication technology.

I/O journal scalability enhancements

The I/O journal is the cornerstone of CDP's ability to provide granular, point-in-time recovery. Every write operation intercepted by the CDP I/O filter or volume filter driver is recorded in the journal, which maintains a chronological record of all data changes. This journal enables administrators to roll back to any point within the short-term retention window, to recover a VM or workload to a state measured in seconds or minutes before a failure occurred.

In v13.0.1, the short-term journal retention window has been extended to a maximum of 168 hours (7 days), a substantial increase from the 24-hour ceiling in earlier versions. This sevenfold expansion of the journal window means that organizations now have a full business week of granular restore points available at any time, dramatically improving the ability to recover from events that are not immediately detected, such as data corruption, logic bombs, or stealthy ransomware encryption that may not manifest for several days.

Journal architecture and storage mechanics

The I/O journal operates on a transaction log model. As I/O operations are intercepted, they are written into transaction log files on the target datastore. Each transaction log file contains a sequence of data blocks along with metadata describing which virtual disk sectors were modified and at what timestamp. Short-term restore points are essentially pointers into these transaction log files, identifying the exact set of I/O operations needed to reconstruct a VM's disk state at a given moment.

The journal retention mechanism works with both short-term and long-term restore points. When a short-term restore point expires beyond the configured retention window, Veeam Backup & Replication commits the data from the corresponding transaction log entries into the nearest technical point. If a transaction log file no longer contains data for any active short-term restore points, it is deleted as redundant, its contents having been consolidated into technical delta disk files. This compaction process ensures that journal storage consumption remains predictable and proportional to the rate of change on protected workloads, even with the extended 168-hour window.

Note

The journal's storage footprint is directly proportional to the I/O change rate of the protected workload and the configured short-term retention period. When planning datastore capacity for CDP targets, multiply the average daily change rate by the number of retention days (up to 7) to estimate peak journal storage requirements.

Scalability benchmarks

The CDP engine in Veeam Backup & Replication has been progressively optimized across recent releases. Internal scalability testing demonstrates that a single Veeam Backup & Replication server configured with 16 vCPU cores and 32 GB of RAM can protect up to 2,000 virtual machines (with 3 virtual disks each) at 15-second RPO intervals, representing approximately a threefold improvement over earlier versions. Additionally, stable performance has been validated for up to 250 VMs (approximately 500 virtual disks) per individual target ESXi host.

The following table summarizes the key scalability parameters for the CDP engine in v13.0.1:

Parameter	Value
Maximum VMs per backup server (15s RPO)	~2,000 VMs (3 disks each)
VMs per target ESXi host (stable)	~250 VMs (~500 virtual disks)
Minimum RPO	2 seconds
Short-term journal retention (max)	168 hours (7 days)
Long-term restore points (max)	Configurable (default: 7 days)
Backup server spec (reference)	16 vCPU, 32 GB RAM

Table 8.9: CDP engine scalability parameters in v13.0.1

Note

Best practice:

While these figures represent validated upper bounds, production environments should be sized with headroom for I/O spikes and growth. Monitor CDP proxy CPU and network utilization alongside journal storage consumption to identify bottlenecks before they affect RPO compliance.

Technet24

CDP I/O filter cross-compatibility

One of the operational challenges in previous versions was the tight coupling between the Veeam Backup & Replication version and the CDP I/O filter version deployed on ESXi hosts. Upgrading Veeam Backup & Replication often required a synchronized update of the I/O filter across all clusters, a process that could involve placing hosts into maintenance mode and potentially disrupting production workloads.

Starting with Veeam Backup & Replication v12.2 and continuing into v13.0.1, Veeam introduced CDP I/O filter cross-compatibility. This feature decouples the I/O filter version from the backup server version, allowing the CDP I/O filter to operate correctly with both the current and the immediately preceding major version of Veeam Backup & Replication. In practical terms, this means administrators can upgrade their backup server to v13.0.1 while continuing to run the existing I/O filter on ESXi hosts, then update the filter on a cluster-by-cluster basis during planned maintenance windows.

This cross-compatibility capability offers several operational advantages:

- **Rolling upgrades**: Upgrade the backup server and CDP proxies first, then update I/O filters on ESXi clusters during scheduled maintenance without interrupting CDP replication.
- **Reduced maintenance windows**: Eliminate the requirement for simultaneous, infrastructure-wide upgrades that previously demanded extended change windows.
- **Improved resilience during upgrades**: CDP policies continue to protect workloads throughout the phased upgrade process, ensuring that RPO commitments are maintained.
- **Multi-cluster flexibility**: Organizations with numerous vSphere clusters across multiple sites can stagger I/O filter updates across different locations and schedules.

Note

Important:

While the I/O filter supports cross-compatibility with the previous major version, it is always recommended to update to the latest filter version as soon as operationally feasible to benefit from the latest performance optimizations and bug fixes. Verify the I/O filter status from the vSphere cluster node in the Veeam Backup & Replication console after upgrading.

Universal CDP: Agent-based continuous data protection

Perhaps the most transformative CDP enhancement in Veeam Backup & Replication v13.0.1 is the introduction of Universal CDP—a technology that decouples continuous data protection from the VMware vSphere hypervisor layer and extends it to any supported Windows workload, regardless of whether it runs on a physical server, a virtual machine on any hypervisor, or a cloud instance.

In previous versions, CDP was exclusively available for VMware vSphere virtual machines through the VAIO (vSphere APIs for I/O Filtering) framework. This approach, while highly effective for vSphere environments, inherently limited CDP to workloads running on VMware infrastructure. Universal CDP removes this limitation by introducing an agent-based architecture that intercepts I/O at the operating system level rather than at the hypervisor level.

Agent-based architecture

Universal CDP relies on two components installed directly on the source workload:

- **Veeam CDP agent service**: This service manages the lifecycle of the volume filter driver, receives captured I/O data, aggregates it into efficient transfer units, and transmits it to the designated CDP proxies.
- **Veeam CDP volume filter driver**: A kernel-mode driver that attaches to each protected volume on the source workload. The driver intercepts I/O transactions generated by the operating system, capturing every write operation in real time and forward them to the CDP agent service for processing and transmission.

Technet24

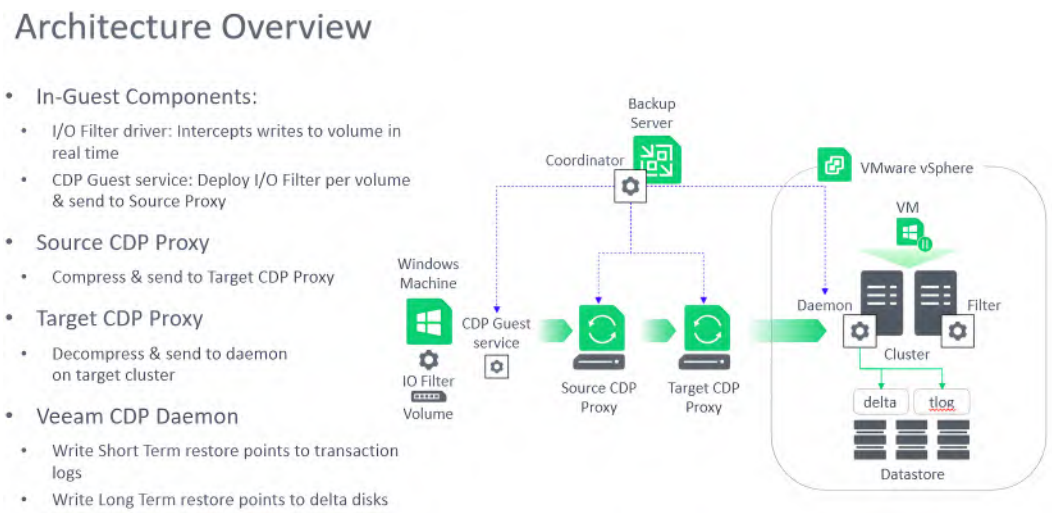


Figure 8.19: Agent-based architecture overview

This agent-based model means that the source workload is entirely abstracted from the underlying infrastructure. A physical server running Windows server, a Hyper-V virtual machine, a cloud instance on Azure or AWS, or even a VMware VM on a non-VAIO-compatible configuration can all be protected with the same near-zero RPO guarantees that were previously exclusive to VAIO-enabled vSphere environments.

Replication targets

Universal CDP replicates source workloads to VMware vSphere clusters or hosts as the target infrastructure. In v13.0.1, supported replication targets include the following:

- **VMware vSphere clusters and standalone hosts**: The Veeam CDP coordinator service communicates with the target vCenter server to create replicas with empty disks based on a synthetic disk layout derived from the source workload's storage configuration.
- **VMware cloud director**: Universal CDP extends to cloud director environments, enabling service providers to offer DRaaS powered by agent-based CDP through existing Veeam cloud connect replication infrastructure.

The target host must have the CDP I/O filter installed and be part of a cluster managed by a vCenter server that is registered in the Veeam Backup and Replication infrastructure. The Veeam CDP coordinator service on the backup server orchestrates the entire replication workflow, querying the source agent for workload metadata, constructing a synthetic disk layout, creating the replica VM on the target, and coordinating the initial synchronization followed by continuous incremental replication.

Universal CDP data flow

The data flow for Universal CDP follows a well-defined pipeline that mirrors the proven CDP for vSphere architecture while accommodating the agent-based source:

1. The Veeam CDP coordinator service queries the Veeam CDP agent service on the source workload for information about the guest OS, hardware configuration, network adapters, and storage layout.
2. The agent service transforms the physical or virtual storage layout into a synthetic disk layout compatible with VMware virtual disk formats.
3. The coordinator service requests the target vCenter server to create a replica VM with empty disks matching the synthetic disk layout.
4. The coordinator service selects appropriate source and target CDP proxies, configures network traffic rules, and initiates the initial full synchronization.
5. Upon completion of the initial sync, the CDP volume filter driver begins intercepting all write operations on the source workload and forwarding them through the agent service to the source proxy.
6. The source proxy compresses and optionally encrypts the data, then transmits it to the target proxy.
7. The target proxy writes the data to the replica's transaction log files on the target datastore, maintaining the I/O journal for short-term restore point creation.

That covers the agent CDP protection including architecture, targets, and data flow. Now, let us look at some of the new failover and recovery enhancements in Veeam Backup and Replication 13.

Failover and recovery enhancements

Veeam Backup & Replication 13 has added some enhancements for failover and recovery, which we will discuss in the following sections.

I/O throughput visualization during failover

When initiating a failover to a CDP replica, administrators must select a restore point from a point-in-time selection timeline. In v13.0.1, the restore point selection wizard now highlights periods of significant I/O throughput directly on the point-in-time selection ribbon. This visual indicator makes it immediately apparent when a large volume of changes occurred within the protected VM, whether due to legitimate bulk operations, database maintenance, or potential ransomware activity that generated a spike in write operations.

This enhancement provides a critical forensic capability during incident response. When recovering from a suspected ransomware attack, administrators can visually identify the

Technet24

moment when encryption activity began, which is typically characterized by a sudden, sustained increase in I/O throughput, and select a restore point immediately before that threshold, ensuring that the recovered workload is free of encrypted or corrupted data.

Planned CDP failover

In addition to the existing support for partial failovers and full site failovers, Veeam Backup & Replication v13.0.1 introduces planned CDP failover. This capability enables a controlled, graceful switchover between production VMs and their CDP replicas with minimal or no interruption to application availability. Planned failover is designed for the following scenarios:

- **Scheduled infrastructure maintenance**: Organizations can migrate workloads to the DR site before performing maintenance on the production infrastructure, ensuring that mission-critical applications remain available throughout the maintenance window.
- **Proactive disaster avoidance**: When advance warning of a natural disaster (such as an approaching hurricane or flood threat) is received, organizations can preemptively switch operations to the DR site before the event impacts the production location.
- **Cloud migration and site relocation**: For service providers, planned failover provides a data-loss-free migration path to move customer workloads from on-premises infrastructure to IaaS platforms running vSphere or VMware cloud director.

After the planned failover is completed and the original production environment is restored, administrators can use the standard failback process to transfer all accumulated changes from the replica back to the production VMs, resuming normal CDP replication.

File-level restore from CDP replicas

Veeam Backup and Replication v13.0.1 also introduces file-level restore functionality for CDP replicas. Previously, recovery from a CDP replica required a full VM failover, even when the actual need was to recover a single file or a small set of files. With file-level restore, administrators can mount the replica's disk state at a specific point in time and extract individual files or folders without initiating a full failover operation.

This capability is particularly valuable in scenarios where an application configuration file has been accidentally modified, a database file has been corrupted, or a user has deleted critical documents. Rather than failing over an entire VM to recover a single file, administrators can surgically extract what is needed to reduce recovery time and minimize disruption to the production environment.

CDP with Veeam cloud connect for DRaaS

Veeam cloud and service providers (VCSPs) benefit significantly from the CDP enhancements in v13.0.1. Both CDP for VMware vSphere and Universal CDP are fully supported within the

Veeam Cloud Connect replication framework, enabling service providers to deliver differentiated **disaster recovery as a service** (**DRaaS**) offerings that leverage near-zero RPO protection.

The CDP infrastructure in a cloud connect environment follows a multi-tenant architecture where components are distributed between the tenant side and the service provider side, communicating through the cloud gateway. The scalability improvements translate directly to service provider economics—the ability to protect up to 2,000 VMs per backup server and 250 VMs per target host means providers can achieve greater tenant density on fewer resources, improving the cost efficiency of their DRaaS platform.

Universal CDP adds a particularly compelling dimension for service providers, as it enables protection of tenant workloads that run outside of VMware environments, including physical servers, Hyper-V VMs, or cloud instances, and broadens the addressable market for CDP-based DRaaS beyond the traditional vSphere customer base.

Design considerations

When implementing CDP in a v13.0.1 environment, architects and administrators should account for the following design factors:

- **Journal sizing**: The extended 168-hour journal retention window requires careful capacity planning on target datastores. Journal storage consumption is directly proportional to the I/O write rate of protected workloads multiplied by the retention period. For write-intensive workloads such as database servers or OLTP systems, journal storage can grow substantially over a 7-day window.
- **Network bandwidth**: CDP continuously transmits I/O data from source to target. Ensure that WAN links between production and DR sites are sized to accommodate the sustained write throughput of all CDP-protected workloads, plus headroom for traffic spikes. The compression and deduplication performed by CDP proxies reduces bandwidth consumption, but network capacity remains a critical sizing parameter.
- **CDP proxy placement**: Deploy at least one source proxy to the production site and one target proxy at the DR site. For Universal CDP workloads, ensure that the source proxy can communicate with all protected machines and that the agent service has been installed and a post-installation reboot completed on each source workload.
- **Coordinator service availability**: The Veeam CDP coordinator service is essential for configuration management, proxy selection, and long-term restore point creation. If the coordinator becomes unavailable, existing CDP policies continue to create and remove short-term restore points, but long-term restore points will not be generated.

Technet24

In environments using the VSA High Availability cluster, coordinator resilience is inherently addressed by the active/passive failover architecture.

- **Universal CDP agent deployment**: The CDP agent and volume filter driver are deployed through protection groups in the Veeam Backup & Replication console. After agent installation, a reboot of the source workload is required to activate the volume filter driver. Plan agent deployments during maintenance windows and verify that the CDP agent service is running and the driver is attached before creating CDP policies.

Tip

Use Veeam ONE monitoring with CDP-specific alarms to track RPO compliance, journal storage consumption, and proxy performance across your CDP infrastructure. The integration of Veeam ONE reports directly into the Veeam Backup & Replication console in v13.0.1 makes it straightforward to maintain visibility over CDP health without switching between management interfaces.

Summary

This chapter has provided you with an overview of continuous data protection (CDP) including Universal CDP within Veeam Backup & Replication v13, as well as the components that comprise the installation. We discussed what is new (including Universal CDP) in the new release. Then, we addressed how to configure and improve network resilience with CDP. Lastly, we examined the enhanced I/O journal scalability features and what else is new for version 13.

This chapter has given you a good understanding of the features of CDP, including the new Universal CDP. The next chapter will cover some more new enhancements within Veeam Backup & Replication v13, like Linux-only backup infrastructure and automatic updates.

Further reading

- CDP Proxies - `https://helpcenter.veeam.com/docs/vbr/userguide/cdp_proxy.html?ver=13`
- Adding CDP Proxy - Wizard - `https://helpcenter.veeam.com/docs/vbr/userguide/cdp_proxy_add.html?ver=13`
- CD Proxy Best Practices - `https://bp.veeam.com/vbr/Support/configurations/cdp_proxy.html`
- Installing CDP Agent Service and Filter Driver - `https://helpcenter.veeam.com/docs/vbr/userguide/uni_cdp_service_install.html?ver=13`

Part 3

Enhancements, Backup Infrastructure, Veeam ONE v13 & Recovery Options

The objective of this section is to teach the reader about the many new enhancements that Veeam has put into this product release. This section will also include many new changes to the backup infrastructure, like PostgreSQL 17.6. We look at the all new Veeam ONE v13 and the many enhancements for monitoring & reporting. We also look at the many recovery options now offered for disaster recovery of your environment.

This part of the book includes the following chapters:

- *Chapter 9, Other Enhancements – Linux Only Infrastructure and Automatic Upgrades*
- *Chapter 10, Backup Infrastructure – PostgreSQL 17.6, Improved Job Manager Scalability, and Other Changes*
- *Chapter 11, Veeam ONE v13*
- *Chapter 12, Recovery Options*

Technet24

9

Other Enhancements – Linux Only Infrastructure and Automatic Upgrades

Veeam Backup & Replication v13 is part of the **Veeam Data Platform**, the newest release of Veeam and has many new enhancements which we will explore in this chapter. We will discuss the new 64-bit Veeam architecture and how this enhances the overall infrastructure of the product for the backup administrator. You will learn how to deploy Linux-only backup infrastructure now with the release of version 13. We will also look at the automatic upgrade implementation and process, along with several other enhancements. These changes will help the backup administrator manage the daily backup environment and help with the security of the product to be aligned with the security department.

In this chapter, we're going to cover the following main topics:

- The new 64-bit architecture
- Linux-only backup infrastructure
- Centralized updates with Veeam updater
- Custom role-based access control (RBAC)

The new 64-bit architecture

Veeam Backup & Replication v13.0.1 represents one of the most architecturally significant releases in the product's history. At the foundation of these changes lies a complete transition to a 64-bit-only platform—a shift that touches every layer of the backup infrastructure, from the backup server itself to the proxies, repositories, and agents that comprise the data

Technet24

protection pipeline. This is not merely a re-compilation of existing code for a wider address space; it is a deliberate, ground-up modernization of the communication protocols, hashing algorithms, data movers, and authentication mechanisms that define how Veeam components interact with one another and with the workloads they protect.

For administrators who have managed Veeam environments across multiple major releases, the implications of this architectural shift are substantial. The removal of 32-bit operating system support across all features and functionality, the replacement of legacy Microsoft RPC and WMI protocols with cross-platform gRPC, and the adoption of the BLAKE3 hashing algorithm collectively deliver a platform that is faster, more secure, and better positioned for the Linux-native future that the Veeam software appliance (VSA) represents. Understanding these changes is essential not only for planning a successful upgrade but also for realizing the full performance and security benefits that v13.0.1 brings to the table.

64-bit-only: The end of 32-bit support

Beginning with v13, Veeam has globally dropped 32-bit operating system support across all features and functionality. This is a clean break—there is no compatibility shim, no legacy mode, and no workaround for environments that still rely on 32-bit infrastructure components. The decision reflects both the realities of the modern enterprise landscape, where 64-bit operating systems have been the standard for well over a decade, and Veeam's strategic need to eliminate the complexity of maintaining parallel 32-bit and 64-bit code paths.

The practical impact of this change varies depending on the role a given server plays in the backup infrastructure. The backup server itself, all backup infrastructure components (proxies, repositories, gateways, WAN accelerators), and the configuration database server must all run 64-bit operating systems. On the protected infrastructure side, application-aware processing for virtual machines is now supported exclusively for 64-bit guest operating systems. It is important to note that crash-consistent, host-level VM backup remains possible regardless of the guest OS architecture. This is a critical distinction for organizations that still have legacy 32-bit workloads in their environment.

Minimum supported operating systems

The 64-bit mandate is accompanied by significant updates to the minimum supported versions across the entire platform. *Table 9.1* summarizes the key changes for backup infrastructure components and protected workloads:

Component / Platform	Dropped support	Minimum version (v13)
Windows server (Infrastructure)	Server 2012, 2012 R2	Windows server 2016

Component / Platform	Dropped support	Minimum version (v13)
Windows server (Protected VMs)	Server 2008, 2012	Windows server 2012 R2
Windows client	Windows 7, 8.x	Windows 10 (1507 LTSC / 22H2 GA)
VMware vSphere	vSphere 6.x	vSphere 7.0 (ESXi and vCenter)
Microsoft Hyper-V	2012, 2012 R2	Hyper-V 2016
SQL server (Config DB)	SQL Server 2012, 2014	SQL server 2016
CentOS (Infrastructure)	All versions	Support dropped entirely
RHEL (Infrastructure)	RHEL 7, 8.0–8.5	RHEL 8.6

Table 9.1: Minimum supported platform versions in Veeam Backup & Replication v13.0.1

Note

Important:

Before upgrading to v13.0.1, verify that every server in your backup infrastructure (including remote proxies, repository servers, tape servers, and WAN accelerators) is running a 64-bit operating system that meets the minimum version requirements. The upgrade wizard will check for compatibility, but discovering issues at that stage can delay your maintenance window.

gRPC: Replacing legacy communication protocols

One of the most consequential architectural changes in v13.0.1 is the complete elimination of Microsoft remote procedure call (RPC) and windows management instrumentation (WMI) as the communication layer between backup infrastructure components and protected workloads. In their place, Veeam has adopted **Google remote procedure call** (**gRPC**), which is a high-performance, open-source remote procedure call framework.

Microsoft RPC and WMI have been the backbone of Windows-based systems management for decades, but they carry significant baggage in a modern security context. RPC relies on dynamic port allocation across a wide range (typically TCP 49152–65535), creating substantial firewall management complexity in segmented networks. WMI has been repeatedly exploited

Technet24

as a lateral movement vector in ransomware and advanced persistent threat (APT) campaigns. Both protocols are inherently Windows-centric, making them incompatible with Veeam's strategic direction toward a cross-platform, Linux-native architecture.

The transition to gRPC addresses all three of these concerns simultaneously. gRPC operates over HTTP/2, uses well-defined ports, and supports TLS encryption natively. It is platform-agnostic by design, enabling the same communication protocol to function identically, whether the backup server runs on Windows or on the Rocky Linux-based VSA. This protocol unification is a key enabler of Veeam's vision for a truly cross-platform backup infrastructure.

Impact on network architecture and firewall rules

The shift to gRPC significantly simplifies network port requirements. Instead of managing wide dynamic port ranges for RPC, administrators can now configure precise, predictable firewall rules. The v13.0.1 REST API and Web UI operate on port 443 (HTTPS), and gRPC communication uses defined, static ports that are documented in the Veeam help center. This reduction in network surface area is particularly valuable for organizations operating under strict micro-segmentation policies or in environments governed by compliance frameworks such as PCI DSS, HIPAA, or NIST 800-171.

> **Note**
>
> **Planning note:**
>
> The transition from RPC/WMI to gRPC requires firewall rule updates across your backup infrastructure. Coordinate with your network security team during the upgrade planning phase. Port 443 must be available on the backup server for the REST API and Web UI. If another application (such as IIS) is already bound to port 443, the Veeam services will fail to start. Resolve port conflicts before initiating the upgrade.

The following table looks at the changes in communication protocols between Veeam Backup and Replication v12 and the newest v13.

Characteristic	v12.x (RPC/WMI)	v13.0.1 (gRPC)
Transport protocol	Microsoft RPC over TCP	HTTP/2 with TLS
Port requirements	Dynamic range (49152–65535) plus static ports	Static, well-defined ports
Platform support	Windows-only	Cross-platform (Windows and Linux)

Characteristic	v12.x (RPC/WMI)	v13.0.1 (gRPC)
Encryption	Optional (RPC encryption)	TLS by default
Firewall complexity	High (dynamic port allocation)	Low (static, predictable)
Security exposure	Known attack surface (WMI lateral movement)	Minimal (modern protocol, encrypted)

Table 9.2: Communication protocol comparison between v12.x and v13.0.1

Now that we have investigated the communication protocol changes, let's look at the authentication changes that have been introduced in v13.

Authentication modernization: NTLM deprecation and certificate-based authentication

Alongside the protocol changes, v13.0.1 deprecates NTLM authentication in favor of Kerberos for domain-joined environments and introduces the **Veeam deployment kit** for certificate-based authentication in workgroup or cross-domain scenarios. NTLM has long been recognized as a weak authentication protocol, vulnerable to pass-the-hash attacks, relay attacks, and credential theft. Its removal from the Veeam communication stack significantly reduces the attack surface of the backup infrastructure.

The Veeam deployment kit

The **Veeam deployment kit** is a new mechanism in v13.0.1 that enables certificate-based authentication between the backup server and infrastructure components. When installed on a target server, the deployment kit allows Veeam Backup & Replication to authenticate using certificates rather than traditional username and password credentials. This is particularly useful for environments that do not support NTLM or Kerberos, which is a common scenario when the backup server runs on the Linux-based VSA and needs to manage Windows infrastructure components.

The deployment kit can be created and exported from either the Veeam console or the Web UI. The exported package is then installed on the target Windows server, after which Veeam can connect using the **"Connect using certificate-based authentication"** option during infrastructure component registration. For environments deploying the JeOS-based Veeam infrastructure appliance (VIA) for proxy and repository roles, certificate-based authentication is enabled by default—no additional configuration is required.

The deployment kit also eliminates the requirement for `admin$` administrative share on protected Windows machines after the initial Veeam agent deployment. This provides a

Technet24

significant security improvement, as the `admin$` share has historically been a target for lateral movement in ransomware campaigns.

BLAKE3 hashing: Faster, more secure data integrity

Veeam v13.0.1 replaces the legacy MD5 hashing algorithm with **BLAKE3**, a modern cryptographic hash function that delivers improvements in both performance and security. BLAKE3 is designed for parallelism and leverages single instruction, multiple data (SIMD) instructions available in modern 64-bit processors, allowing it to process data at speeds that approach the memory bandwidth of the system rather than being constrained by CPU throughput.

In practical terms, the adoption of BLAKE3 reduces backup proxy and backup agent CPU usage by up to 30 percent in scenarios where the CPU is the bottleneck. This is not an incremental improvement but a step change that directly translates into increased backup throughput without requiring additional hardware. For organizations running large-scale environments with hundreds of concurrent backup streams, the cumulative effect of this CPU reduction can be the equivalent of adding several additional proxy servers to the infrastructure.

Beyond performance, BLAKE3 provides stronger security guarantees than MD5. While MD5 has been cryptographically broken for over a decade (with practical collision attacks demonstrated as far back as 2004), BLAKE3 offers strong resistance against both collision and pre-image attacks. In the context of backup data integrity verification, particularly during restore operations where you need absolute confidence that backup data has not been tampered with, this cryptographic strength is essential.

The following table outlines the different hashing algorithms that have changed between Veeam Backup and Replication v12 and v13.

Property	**MD5 (v12.x)**	**BLAKE3 (v13.0.1)**
Hash output size	128-bit	256-bit (default)
Collision resistance	Broken (practical attacks since 2004)	Strong (no known weaknesses)
Parallelism	Sequential processing	SIMD-optimized, tree-based parallelism
CPU impact on proxies	Baseline	Up to 30% reduction

Property	MD5 (v12.x)	BLAKE3 (v13.0.1)
64-bit optimization	Minimal	Full AVX2/AVX-512 utilization

Table 9.3: Hashing algorithm comparison between v12.x and v13.0.1

Instant recovery engine and data mover improvements

The 64-bit architecture enables substantial improvements to the instant recovery engine, which is the core technology behind Veeam's ability to run virtual machines directly from backup files. In v13.0.1, the instant recovery engine delivers up to 50 percent improved I/O performance for VMs running from backups compared to v12.x. This improvement directly accelerates not only instant VM recovery scenarios but also all capabilities that rely on this engine, including SureBackup verification, instant disk recovery, and staged restore operations.

The mass VM restore engine has also been significantly optimized, with reduced CPU and RAM consumption on the backup server during large-scale restore operations. Veeam reports that the entire VM restore engine improvements enable up to 1,000 concurrent VM restores from a single backup server, a capability that has obvious implications for disaster recovery scenarios where rapid, parallel recovery of large VM fleets is required.

Agent-based backup performance has seen equally dramatic improvements. The Veeam agent for Microsoft Windows now delivers more than double the backup throughput of v12.x on identical hardware, which is an improvement that reflects the benefits of the rewritten 64-bit data mover engine and the transition to gRPC communication. The same improvements apply to the Veeam agent for Linux. For unstructured data backup workloads, v13.0.1 achieves an average 25 percent performance improvement over v12.x, with additional gains for NFS3 shares using multiple TCP connections.

PostgreSQL 17.6: A unified configuration database

The 64-bit architecture transition is accompanied by an upgrade to PostgreSQL 17.6 as the configuration database engine. This is the same database version used by VSA, ensuring parity between Windows-based and Linux-based deployments. For organizations upgrading from v12.x, the PostgreSQL database is automatically upgraded during the installation process, and previous versions (such as PostgreSQL 15) can be safely removed after confirming stable operation.

PostgreSQL 17.6 brings its own set of performance and reliability improvements, including enhanced query planning, improved vacuum operations, and better handling of large configuration databases that are common in enterprise-scale Veeam deployments with thousands of managed objects. The use of a consistent database engine across both

Technet24

deployment models also simplifies future migration paths for organizations that plan to transition from the Windows-based installation to the VSA.

A cross-platform architecture for the future

Taken together, the architectural changes in v13.0.1 represent a deliberate strategy to create a unified, cross-platform codebase. The adoption of gRPC, the transition to BLAKE3, the PostgreSQL database standardization, and the certificate-based authentication model are the same foundations used by the Veeam Software Appliance. For organizations running Veeam on Windows, the v13.0.1 architecture effectively functions as a steppingstone: upgrade to v13.0.1, confirm stable operation across your backup infrastructure, and gain optionality for a future migration to the appliance model without being forced into it today.

The shared cross-platform codebase also means that performance and scalability improvements apply equally to both Windows and Linux deployments. There is no "*first-class*" platform for v13.0.1. The same data mover engine, the same hashing algorithm, the same communication protocol, and the same authentication mechanisms operate identically regardless of the underlying operating system. This architectural consistency is foundational to Veeam's ability to deliver a truly hybrid backup infrastructure where Windows and Linux components can coexist and interoperate seamlessly.

With Linux-capable mount servers, gateway servers, and guest interaction proxies now available in v13.0.1, organizations can begin building entirely Linux-based backup infrastructures while retaining the option to protect Windows workloads. The 64-bit architecture is the technical enabler that makes this possible.

Now that we have looked at all the changes with the 64-bit architecture, communication protocols, and hashing algorithms, let's now dive into using a Linux-only backup infrastructure with the new v13.

Linux-only backup infrastructure

One of the most transformative changes introduced in Veeam Backup & Replication v13 is the ability to deploy and operate an entirely Linux-based backup infrastructure. For the first time in the product's history, every infrastructure role, from the backup server itself to proxies, repositories, mount servers, gateway servers, tape servers, and guest interaction proxies, can run natively on Linux without any dependency on Microsoft Windows. This represents a fundamental architectural shift that has been years in the making and carries significant implications for security, licensing costs, and operational simplicity. Security is enhanced, licensing costs are reduced (no more Microsoft licensing required), and there is a focus on one operating system (Linux) versus more than one.

While previous versions of Veeam Backup & Replication introduced Linux support for select roles such as backup proxies and hardened repositories, v13 completes the journey by bringing the remaining Windows-only components into the Linux ecosystem. Organizations now have the choice to build their entire data protection platform on a single operating system foundation, dramatically reducing the attack surface and eliminating the need for Windows server licensing across the backup infrastructure.

To appreciate the significance of v13's Linux-only capability, it helps to understand how Veeam's infrastructure role support has evolved across recent major releases. Each version has progressively expanded Linux support, moving from a handful of peripheral roles to full coverage of the entire infrastructure stack.

Version	Linux roles supported	Windows-only roles
v11	Backup proxy (VMware), Linux repository, hardened repository (experimental)	Backup server, mount server, gateway server, guest interaction proxy, WAN accelerator, distribution server
v12	Backup proxy (VMware), Linux repository, hardened repository (GA), CDP proxy	Backup server, mount server, gateway server, guest interaction proxy, WAN accelerator, distribution server
v13	All roles: backup server (VSA), backup proxy, repository, hardened repository, mount server, gateway server, guest interaction proxy, WAN accelerator, log shipping server, distribution server, vPower NFS	None required (Windows console still needed for full feature parity)

Table 9.4: Version comparison for supported roles on Linux and Windows

As *Table 9.4* illustrates, v13 effectively inverts the historical paradigm. Linux is now the primary platform for Veeam infrastructure, with Windows remaining as a supported alternative rather than a requirement. This shift was driven by several factors: community demand from Linux-first organizations, the security advantages of a minimized operating system footprint, and the operational benefits of appliance-based deployment models.

Technet24

New Linux infrastructure roles in v13

While backup proxies and repositories have supported Linux for several versions, v13 introduces Linux support for a collection of infrastructure roles that were previously exclusive to Windows. Each of these newly supported roles eliminates one more reason an organization might need to maintain Windows servers in the backup infrastructure.

Mount server

The **mount server** is a critical infrastructure component required for restore operations that work with guest OS files, application items, and instant recovery workflows. In previous versions, this role could only be assigned to a Windows server, which meant that even organizations running Linux proxies and repositories still needed at least one Windows machine in their environment.

In v13, the mount server role can be assigned to any Linux server in the backup infrastructure, including a Veeam infrastructure appliance (VIA) deployed from the JeOS ISO. When deployed on a VIA, the mount server uses certificate-based authentication over secure communication protocols, eliminating the need for open SSH ports and minimizing the potential attack surface. The Linux mount server supports file-level recovery, instant recovery to VMware vSphere, and application item recovery through the Veeam explorers.

> **Note**
>
> A Linux mount server deployed from the VIA image supports Rocky Linux 9 and RHEL 9. If deploying a mount server on a manually configured Linux host, you must install the `aspnetcore-runtime` (version 8.0 or later), `bindfs`, `samba-client`, and `ntfs-3g` packages before adding the server to the backup infrastructure.

Gateway server

Gateway servers act as intermediaries between the backup server and backup repositories that cannot host Veeam data movers directly, such as object storage targets, deduplication appliances (Dell data domain, HPE StoreOnce), SMB and NFS shares, and NDMP-enabled storage. The gateway server hosts the target data mover and facilitates data transfer between the source proxy and the repository.

With v13, gateway servers can now run on Linux, allowing administrators to place these intermediary components on the same operating system as the rest of the infrastructure. This is particularly beneficial for organizations using object storage as a primary or secondary backup target, as the gateway server handling data flow to cloud or on-premises S3-compatible storage no longer requires a Windows intermediary.

Guest interaction proxy

The **guest interaction proxy** (**GIP**) is responsible for application-aware processing during backup operations. It connects to the guest operating system of virtual machines to coordinate VSS quiescing for Windows VMs, perform pre-freeze and post-thaw scripts for Linux VMs, and handle transaction log truncation for applications such as Microsoft SQL server and exchange.

Moving the GIP role to Linux means that application-aware backup processing no longer forces a Windows dependency into the infrastructure. The Linux-based GIP uses the same VIX or network-based connectivity methods to reach guest operating systems and supports all the same application-aware processing features as its Windows counterpart except for using a gMSA account as the Guest OS credentials.

Additional roles

Beyond the three major roles detailed above, v13 also brings Linux support to several other infrastructure components. The WAN accelerator, used for bandwidth optimization in remote office and disaster recovery scenarios, can now run on Linux. The log shipping server, which handles transaction log backup shipping for database protection, has gained Linux support. The distribution server role for Veeam Agent deployment and the vPower NFS service for Instant Recovery also now operate natively on Linux.

Veeam infrastructure appliance

To simplify the deployment and management of Linux-based infrastructure components, Veeam has introduced the **Veeam Infrastructure Appliance** (**VIA**). The VIA is a purpose-built, pre-hardened appliance image based on Veeam's JeOS, which is itself built on Rocky Linux 9. It is delivered as a bootable ISO image that can be installed on physical servers, virtual machines on any hypervisor, or cloud instances.

The VIA is distinct from the VSA, which serves as the backup server itself. While the VSA hosts the core Veeam Backup & Replication services and the configuration database, VIA images are used to deploy the distributed infrastructure components: proxies, repositories, mount servers, gateway servers, and other roles. Both appliances share the same JeOS foundation and security hardening but serve different purposes in the infrastructure hierarchy.

The following table looks at the differences between using the Veeam VIA to deploy the environment and manual Linux installations.

Technet24

Characteristic	Veeam infrastructure appliance	Manual Linux server
Operating system	JeOS (Rocky Linux 9, pre-hardened)	RHEL 9, Rocky Linux 9, Ubuntu 20.04/22.04, Debian 11/12, SLES 15
Authentication	Certificate-based (automatic)	SSH credentials or certificate-based (manual setup via Deployment Kit)
Security hardening	Pre-hardened, DISA STIG-compliant, SSH disabled by default	Administrator responsibility to harden
OS patching	Managed by Veeam through a dedicated update repository	Administrator responsibility
Deployment options	Two variants: standard and with iSCSI/NVMe support	Full OS installation with manual Veeam component deployment
Veeam support scope	Full stack (OS + Veeam components)	Veeam components only

Table 9.5: Comparison of VIA and manually installed Linux

The VIA comes in two installation variants. The standard VIA includes only the minimal software configuration needed for general-purpose proxy, repository, mount server, VHR (Veeam Hardened Repository) and gateway roles. The second variant, the VIA with iSCSI and NVMe support, adds the additional drivers and tools required for direct SAN transport mode when connecting to iSCSI or NVMe-oF storage targets. From a security perspective, the standard variant is recommended due to requiring less configuration and setup, unless SAN transport is specifically required.

Tip

When deploying infrastructure components, the VIA is the recommended approach over manual Linux server configuration. The VIA provides certificate-based

authentication by default, receives managed OS updates from Veeam, and is fully supported at the operating system level by Veeam technical support.

Certificate-based authentication and reduced port requirements

A significant security enhancement that accompanies the Linux-only infrastructure model is the shift toward certificate-based authentication for infrastructure component communication. When infrastructure components are deployed using VIA, certificate-based authentication is enabled automatically during the enrollment process. The first backup server that connects to a newly deployed VIA establishes a trust relationship through certificate fingerprint verification, after which all subsequent communication is authenticated using certificates rather than credentials.

For organizations that deploy infrastructure components on manually configured Linux servers rather than VIA images, the Veeam deployment kit provides a mechanism to enable certificate-based authentication. The deployment kit is a package generated from the Veeam Backup & Replication server (either through the console UI or via the `Generate-VBRBackupServerDeployerKit` PowerShell cmdlet) that contains temporary certificates and installation packages. Once installed on the target server, the deployment kit enables the Veeam deployer service, which facilitates certificate-based enrollment without requiring SSH credentials.

Certificate-based authentication delivers several security benefits over traditional credential-based approaches:

- It eliminates the need to store and transmit passwords for infrastructure component connections.
- It reduces the attack surface by allowing SSH to be disabled on infrastructure servers after initial deployment.
- It aligns with zero-trust security principles where every connection is verified cryptographically rather than through shared secrets.

The move to Linux and certificate-based authentication has also enabled Veeam to significantly reduce the number of network ports required for infrastructure communication. With credential-based SSH connectivity, multiple ports needed to be opened for management agent deployment, data mover communication, and various service-to-service interactions. Certificate-based communication consolidates much of this traffic, simplifying firewall rules and reducing the number of potential entry points in the network.

Technet24

Designing a Linux-only backup infrastructure

Deploying a complete Linux-only backup infrastructure requires thoughtful planning to ensure all roles are covered and that the design accounts for the operational differences between Linux and Windows-based deployments. The following considerations should guide the architecture:

Core architecture components

Consider these core architecture components while designing a Linux-only backup infrastructure:

- **Veeam software appliance (VSA)**: VSA serves as the backup server, hosting the Veeam Backup & Replication services, the PostgreSQL 17.6 configuration database, and the web-based management console. It is deployed from either an OVA file (for VMware vSphere environments) or the full ISO image. The VSA requires a minimum of two 256 GB disks, 8 CPU cores, and 16 GB of RAM.
- **Veeam infrastructure appliances (VIA)**: One or more VIA instances should be deployed for the distributed infrastructure roles. A common design pattern is to co-locate the proxy and mount server roles on the same VIA, with separate VIA instances for repositories and gateway servers depending on performance and isolation requirements.
- **Hardened repositories**: For immutable backup storage, the VIA can be deployed as a hardened repository. When installed in this mode, the appliance configures XFS with reflink support for fast cloning, enforces single-use credentials that are discarded after initial enrollment or certificate-based authentication, and prevents any subsequent root-level access to the backup data.

Role placement guidelines

The following table shows the Veeam roles, the recommended deployment method and the placement of the server within the infrastructure to get the best performance.

Infrastructure role	Recommended deployment	Placement guidance
Backup proxy	VIA (standard or with iSCSI/ NVMe)	Close to source hypervisors for network mode, or on SAN fabric for direct SAN mode

Infrastructure role	Recommended deployment	Placement guidance
Backup repository	VIA (standard) on server with local or DAS storage	Co-locate with physical storage; use XFS for reflink (fast clone) support
Mount server	VIA (standard), co-located with proxy or repository	Required for FLR and Instant Recovery; place near restore targets
Gateway server	VIA (standard)	Close to target repository for dedup appliances; close to proxy for source-side dedup
Guest interaction proxy	VIA (standard) or co-located with backup proxy	Network access to guest VMs is required for application-aware processing
WAN accelerator	VIA (standard) with dedicated cache disk	One at source site, one at target site; requires fast local storage for cache

Table 9.6: Veeam roles and placement guidance

Windows console considerations

While v13 enables a fully Linux-based infrastructure for all backup processing roles, one important caveat remains: the Veeam Backup & Replication console—the thick client management application—still requires a Windows machine to run. The new web-based UI provides access to the majority of management operations directly from a browser, but certain advanced configuration tasks and features are currently only available through the Windows console. Organizations pursuing a true zero-Windows strategy should plan for either a single Windows administrative workstation (which need not be a server) or use the web UI exclusively, accepting the current feature gap.

As the web UI matures through subsequent releases, the expectation is that it will eventually achieve full feature parity with the Windows console, at which point the last Windows dependency in the Veeam ecosystem will be truly optional.

Technet24

Note

Important:

The Windows console is used for management only and does not process backup data. It can run on any Windows workstation and does not require Windows server licensing. Running the console does not compromise the security benefits of a Linux-only backup processing infrastructure.

Migration considerations

Organizations running existing Veeam Backup & Replication environments on Windows face the question of how to transition to the new Linux-based architecture. As of v13.0.1, the migration path requires careful planning, as there is no automated in-place upgrade from a Windows-based VBR server to the VSA. The recommended approach involves deploying a new VSA alongside the existing Windows installation and migrating jobs, configuration, and repository references to the new environment.

For distributed infrastructure components, the transition is more straightforward. New VIA-based proxies and repositories can be added to the existing infrastructure and gradually take over workloads from their Windows counterparts. Once all jobs have been redirected to the new Linux components, the legacy Windows servers can be decommissioned. This phased approach minimizes disruption and allows organizations to validate the Linux infrastructure before committing to a full migration.

Veeam has acknowledged the importance of robust migration tooling and has indicated that improved migration utilities will be delivered in subsequent releases. In the meantime, organizations should plan for a parallel deployment strategy and should not attempt to rush the transition, particularly in production environments with complex job configurations and extensive repository data.

Tip

Veeam has extended support for v12 to four years from its release date ending in February 2027 (instead of the usual three) specifically to give organizations adequate time to plan and execute their migration to the v13 Linux-based architecture. This extended support applies to all corresponding agent and application plug-in versions.

Security benefits of a Linux-only infrastructure

The move to a Linux-only backup infrastructure delivers measurable security improvements that extend beyond simply changing the operating system. The combination of a hardened JeOS base, certificate-based authentication, and a minimized software footprint creates a fundamentally different security posture compared to traditional Windows-based deployments.

- **Reduced attack surface**: The JeOS distribution contains only the packages and services required for Veeam operations. There is no desktop environment, no unnecessary network services, and no extraneous software that could introduce vulnerabilities. SSH is disabled by default on VIA deployments, and certificate-based authentication eliminates the need for credential storage on infrastructure servers.
- **Elimination of Windows-specific attack vectors**: A Linux-only infrastructure is not susceptible to Windows-specific malware, Active Directory credential theft attacks, or exploitation of Windows services and protocols. Ransomware campaigns that target Windows environments through SMB vulnerabilities, remote desktop protocol, or PowerShell-based lateral movement have no viable attack path against a properly hardened Linux backup infrastructure.
- **Managed patching and compliance**: VIA and VSA deployments receive operating system updates through Veeam's managed update repository, ensuring that patches are tested for compatibility before deployment. This centralized update mechanism reduces the risk of a failed OS patch breaking backup operations and ensures that the infrastructure remains in a known-good state for compliance and audit purposes. The pre-hardened configuration aligns with DISA STIG benchmarks, supporting compliance with regulatory frameworks such as NIS2, DORA, GDPR, and ISO 27001.
- **Mandatory multi-factor authentication**: The VSA and VIA enforce multi-factor authentication for administrator and security officer accounts during initial configuration. Combined with the security operator role (which must approve high-risk operations such as backup data deletion), the Linux-based infrastructure implements defense-in-depth principles at the authentication layer.

By extending Linux support to every infrastructure role such as mount servers, gateway servers, guest interaction proxies, WAN accelerators, log shipping servers, and distribution servers, Veeam has eliminated the last barriers to a fully Windows-free data protection platform.

The introduction of the VIA provides a standardized, pre-hardened, and Veeam-supported deployment model for these roles, while certificate-based authentication and reduced port requirements deliver tangible security improvements. Organizations planning new

Technet24

deployments should strongly consider the Linux-only architecture as their default, while those with existing Windows-based environments can take advantage of extended v12 support to plan a methodical, phased migration to the new platform.

Centralized updates with Veeam updater

One of the most operationally significant enhancements in Veeam Backup & Replication v13.0.1 is the introduction of the **Veeam updater** service, which provides a centralized, automated mechanism for deploying updates across your entire Veeam infrastructure. Prior to this capability, administrators were required to manually update each remote component individually, which is a time-consuming and error-prone process, particularly in large distributed environments with dozens or hundreds of managed servers, proxies, and repositories.

The Veeam updater fundamentally changes this paradigm by allowing the backup server to orchestrate and push updates to all managed infrastructure components from a single control point. This capability is especially valuable for organizations running VSA deployments, where the underlying JeOS-based Linux appliance and its Veeam components can now be updated in a coordinated fashion without requiring individual SSH sessions or manual intervention on each node.

Architecture and components

The centralized update mechanism relies on several components working together to deliver a seamless upgrade experience. Understanding this architecture is essential for troubleshooting and for planning maintenance windows in production environments.

Veeam updater service: The Veeam updater service runs on the Veeam Backup & Replication server and acts as the central orchestration engine. It maintains an inventory of all managed infrastructure components, tracks their current version levels, and coordinates the distribution of update packages. The service communicates with remote components over secure channels and manages the sequencing of updates to minimize downtime and maintain data protection continuity.

Update distribution model: When a new cumulative patch or update becomes available, the Veeam updater downloads the update package to the backup server. From there, it distributes the package to each managed component according to a configurable deployment schedule. The service supports both immediate deployment and scheduled deployment windows, giving administrators flexibility in how updates are rolled out across the environment.

The following table explains the Veeam component and the update behavior of said component.

Component	Update Behavior
Backup server	Updated first as the central orchestrator; requires a service restart to complete the update process.
Backup proxies	Updated in parallel where possible; active proxy tasks are allowed to complete before the update is applied.
Backup repositories	Updated sequentially to maintain data availability; the Veeam Data Mover service is restarted on each repository after the update.
WAN accelerators	Updated during scheduled windows; cache must be repopulated after certain updates that modify the deduplication engine.
Tape servers	Updated individually; requires that no tape jobs are actively running during the update process.
VSA appliance nodes	Both the Veeam application layer and the underlying JeOS operating system packages can be updated centrally through the VSA management interface.

Table 9.7: Managed infrastructure components and their centralized update behavior in Veeam Backup & Replication

Configuring automatic updates

Enabling centralized updates requires minimal configuration, but understanding the available options ensures that updates are deployed in a manner that aligns with your organization's change management policies and maintenance windows.

Enabling the Veeam updater

The Veeam updater is enabled by default in new installations of Veeam Backup & Replication v13.0.1. For environments upgraded from earlier versions, the service may need to be manually enabled.

To configure the Veeam updater, navigate to the main menu in the Veeam Backup & Replication console, select **Updates**, and then **Update Settings**. Here, you will find automatic update settings that control how the backup server checks for and distributes updates.

Technet24

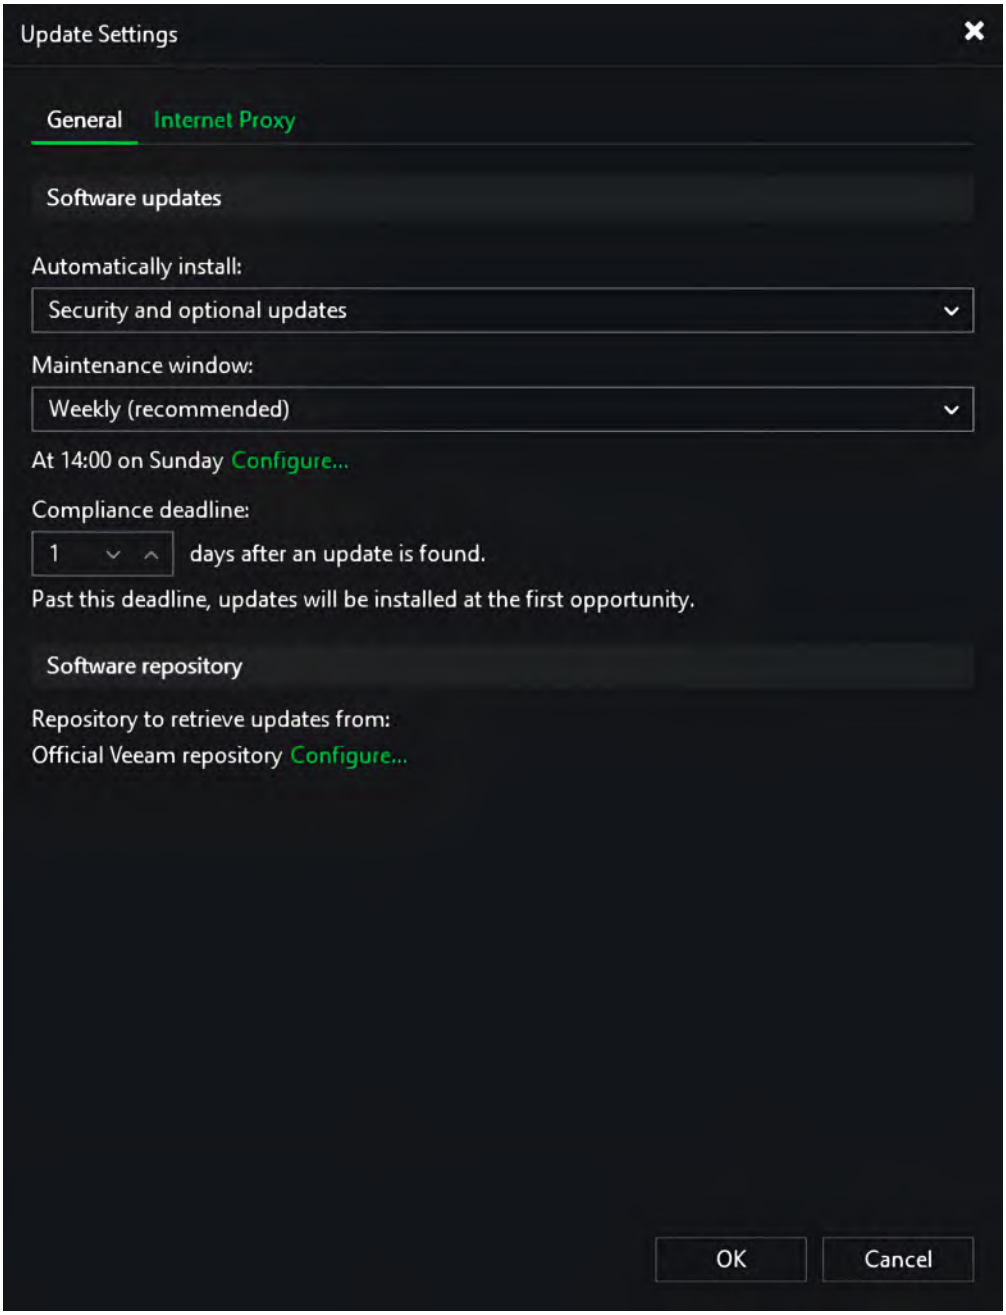


Figure 9.1: Update settings dialog

The configuration options include the ability to enable or disable automatic update checking, configure the frequency of update checks, specify whether updates should be downloaded automatically or require manual approval, and define deployment schedules for pushing updates to managed components.

Update notification and approval workflow

When the Veeam updater detects a new update, it can be configured to handle the update in one of several ways depending on your organizational requirements:

- **Automatic download and install**: The update is downloaded and deployed to all managed components without administrator intervention. This mode is suitable for lab environments or smaller deployments where rapid patching is prioritized.
- **Automatic download with manual approval**: The update package is downloaded automatically, but deployment requires explicit administrator approval. This is the recommended mode for production environments, as it provides a window for testing updates in a staging environment before rolling them out.
- **Manual download and install**: Both download and deployment require manual intervention. This mode provides maximum control and is appropriate for highly regulated environments with strict change management requirements.

Note

Update sequencing:

The Veeam updater enforces a strict update order. The backup server must always be updated before any remote components. This ensures that the central management plane is running the latest version before attempting to update downstream infrastructure. Attempting to update a remote component to a version newer than the backup server will result in a compatibility error.

Scheduling update deployments

For production environments, scheduling updates during defined maintenance windows is critical to minimizing operational disruption. The Veeam updater supports granular scheduling options that allow administrators to specify the days of the week and time windows during which updates may be deployed. This scheduling mechanism respects active job execution. This means that if a backup, replication, or copy job is running on a target component when the scheduled update window opens, the Veeam Updater will wait for the job to complete before proceeding with the update.

Technet24

Additionally, the Veeam updater supports a staged rollout approach. Administrators can create update groups that define which components should be updated first, second, and so on. This is particularly useful in environments with multiple proxy tiers or geographically distributed repositories, where you may want to validate the update on a subset of components before proceeding with the full rollout.

Update process for VSA deployment

Veeam Software Appliance deployments benefit from an additional layer of centralized update management. Because the VSA bundles the Veeam application layer with a hardened JeOS Linux operating system, updates must address both the application components and the underlying OS packages. The v13.0.1 release introduces a unified update mechanism through the VSA management Web UI that coordinates these two update streams.

When an update is available for a VSA node, the management interface displays the available updates along with release notes and any prerequisite conditions. The administrator can review the update details, including whether the update requires a reboot of the appliance, and then initiate the update with a single action. For VSA High Availability clusters, the update process is HA-aware—it will update passive nodes first, perform a controlled failover, and then update the formerly active node, ensuring continuous availability of backup services throughout the update process.

> **Note**
>
> **HA cluster update considerations:**
>
> Before initiating a centralized update on a VSA HA cluster, ensure that all cluster nodes are healthy and that the Patroni service is reporting a normal replication state. Running an update while a cluster node is in a degraded state can result in update failures and potential service disruption.

Monitoring and troubleshooting updates

The Veeam Updater provides comprehensive logging and monitoring capabilities that allow administrators to track the progress of update deployments and diagnose any issues that may arise during the update process.

Update history and status

The update history for all managed components is accessible from the Veeam Backup & Replication console under the **History** node. Each update event is logged with a timestamp, the component name, the update version, and the result status (success, warning, or failure).

For failed updates, the log entry includes a detailed error message and a link to the relevant log files for further investigation.

Common troubleshooting scenarios

Table 9.8 consolidates some common troubleshooting scenarios:

Symptom	Likely cause	Resolution
Update download fails	Firewall blocking outbound HTTPS to Veeam update servers	Verify that the backup server can reach `download.veeam.com` on port 443. Configure proxy settings if required.
Remote component update times out	Network connectivity issues or the Veeam Installer Service is not running on the target	Verify network reachability and confirm that the Veeam Installer Service is running. Check Windows Firewall rules on the target.
Update succeeds but services fail to start	Insufficient disk space or conflicting third-party software	Check available disk space on the system drive. Review the Veeam service logs for specific error messages and address any conflicts.
VSA appliance update fails with Patroni error	HA cluster node is in a degraded state	Resolve the HA cluster health issue before retrying. Use `'patronictl list'` to diagnose the cluster state and `'systemctl restart patroni'` if needed.

Technet24

Symptom	Likely cause	Resolution
Version mismatch after update	Update was partially applied due to an interruption	Re-run the update from the Veeam Updater. If the issue persists, manually download the update package and apply it directly to the affected component.

Table 9.8: Common centralized update troubleshooting scenarios and resolutions

Custom role-based access control (RBAC)

Veeam Backup & Replication v13.0.1 introduces a significantly enhanced RBAC framework that allows administrators to define custom roles with granular permissions tailored to their organization's specific operational requirements. While previous versions of Veeam provided a fixed set of predefined roles (Veeam backup administrator, Veeam backup operator, Veeam restore operator, and Veeam backup viewer), the custom RBAC capability in v13.0.1 allows fine-grained control over which actions users can perform within the backup infrastructure.

This enhancement is particularly important for organizations that need to delegate specific backup management responsibilities to different teams, such as allowing a help desk team to perform restores without granting them access to modify backup job configurations or enabling a storage team to manage repository settings without exposing broader infrastructure configuration options.

Understanding the RBAC architecture

The custom RBAC system in v13.0.1 is built on a permissions-based model that associates specific capabilities with roles, and roles with user or group accounts. This architecture provides three layers of access control that work together to define what any given user can see and do within the Veeam console.

Permissions, roles, and assignments

Permissions represent individual actions or capabilities within the Veeam environment, such as "*Create Backup Job*," "*Perform Instant Recovery*," or "*Manage Repository Settings*." Veeam v13.0.1 exposes over 80 discrete permissions that can be combined into custom roles. Each permission is categorized into a functional area (such as backup, restore, infrastructure, or configuration) to simplify the process of constructing roles.

Roles are named collections of permissions that represent a logical job function or responsibility within the organization. Roles can be either built-in (the predefined roles from

earlier versions) or custom-defined by administrators. Custom roles can include any combination of the available permissions, allowing for precise alignment with organizational responsibilities.

Assignments link roles to specific Active Directory users or groups and, optionally, to specific infrastructure objects. This means that a role can be scoped not only by what actions it permits, but also by which infrastructure objects (servers, repositories, jobs) those actions can be performed on. For example, a custom "*Branch Office Operator*" role might grant restore permissions only for backup jobs associated with a specific branch office.

Note

Active Directory integration:

Custom RBAC roles are assigned to Active Directory users and groups. This means that the Veeam Backup & Replication server must be joined to an Active Directory domain (or have a trust relationship with the domain containing the target accounts) for custom RBAC to function. Local Windows accounts are also supported for the full scope-based assignment capabilities of custom roles.

Built-in roles vs. custom roles

Before creating custom roles, it is important to understand the capabilities and limitations of the built-in roles, as they serve as both a starting point for customization and as a fallback for common access patterns:

Built-in role	Capabilities	Limitations
Backup administrator	Full access to all Veeam functions, including configuration, job management, restore operations, and infrastructure management.	Cannot be scoped to specific objects; always has global access.
Backup operator	Can start, stop, and retry backup and replication jobs. Can perform basic monitoring and reporting.	Cannot create, modify, or delete jobs. Cannot manage infrastructure.
Restore operator	Can perform restore operations including full VM restore, file-level recovery, and application item recovery.	Cannot manage backup jobs or infrastructure. Limited to restore actions only.

Technet24

Built-in role	Capabilities	Limitations
Backup viewer	Read-only access to the Veeam console. Can view job status, session history, and infrastructure inventory.	Cannot perform any actions; strictly a monitoring and audit role.

Table 9.9: Built-in RBAC roles in Veeam Backup & Replication and their capabilities

Custom roles address the gaps between these built-in roles. For example, if your organization needs a role that can both operate backup jobs and perform restores (combining elements of backup operator and restore operator) but should not have access to tape infrastructure, a custom role can be created that includes precisely those permissions.

Creating custom roles

Custom roles are created and managed through the Veeam Backup & Replication console under the **Users and Roles** section, accessible from the main menu. The role creation wizard guides administrators through the process of selecting permissions, naming the role, and optionally defining a scope.

1. Open the Veeam Backup & Replication console and navigate to the main menu. Select **Users and Roles** to open the security configuration interface.

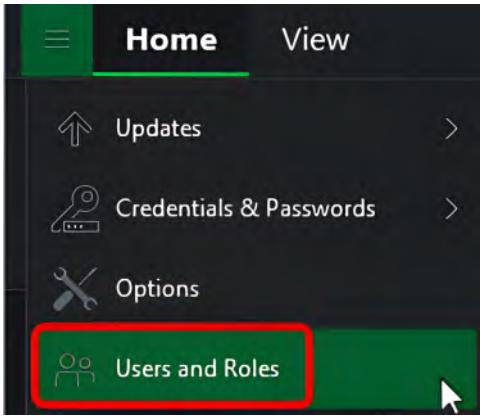


Figure 9.2: Users and Roles selection in main menu

2. Click the **Roles** tab to view existing roles, then click **Add...** to create a new custom role.

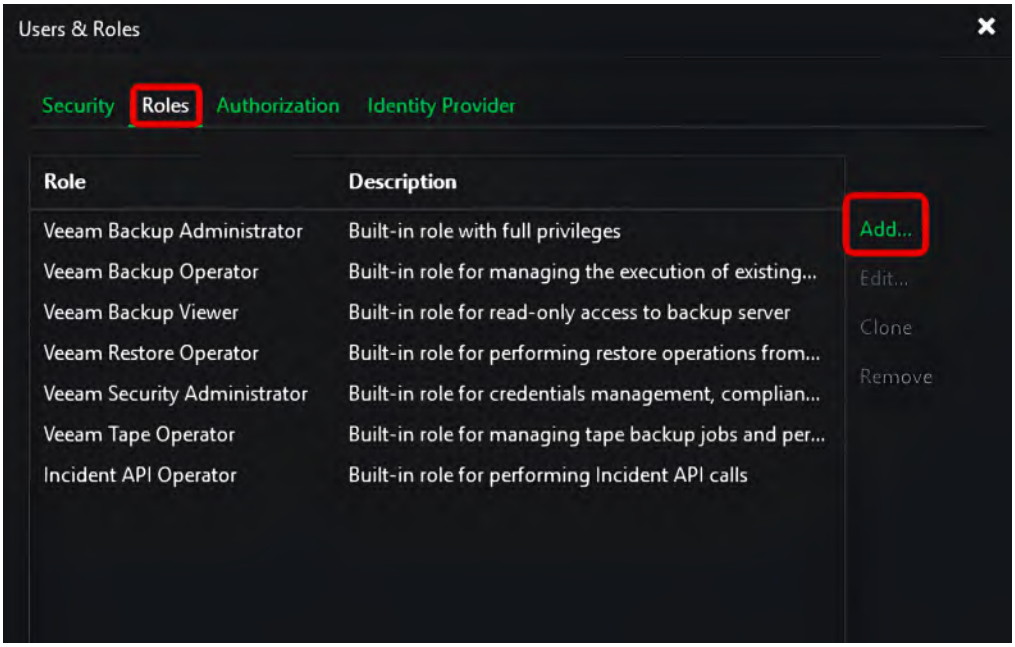


Figure 9.3: Roles tab in the Users and Roles dialog

3. Provide a descriptive name for the role (for example, "*Branch Office Restore Operator*" or "*Storage Team – Repository Manager*"). Include a description that documents the role's intended purpose and target audience. Select the checkbox for **Manage Backups** or **Manage restores** or both to see the other required tabs.

Technet24

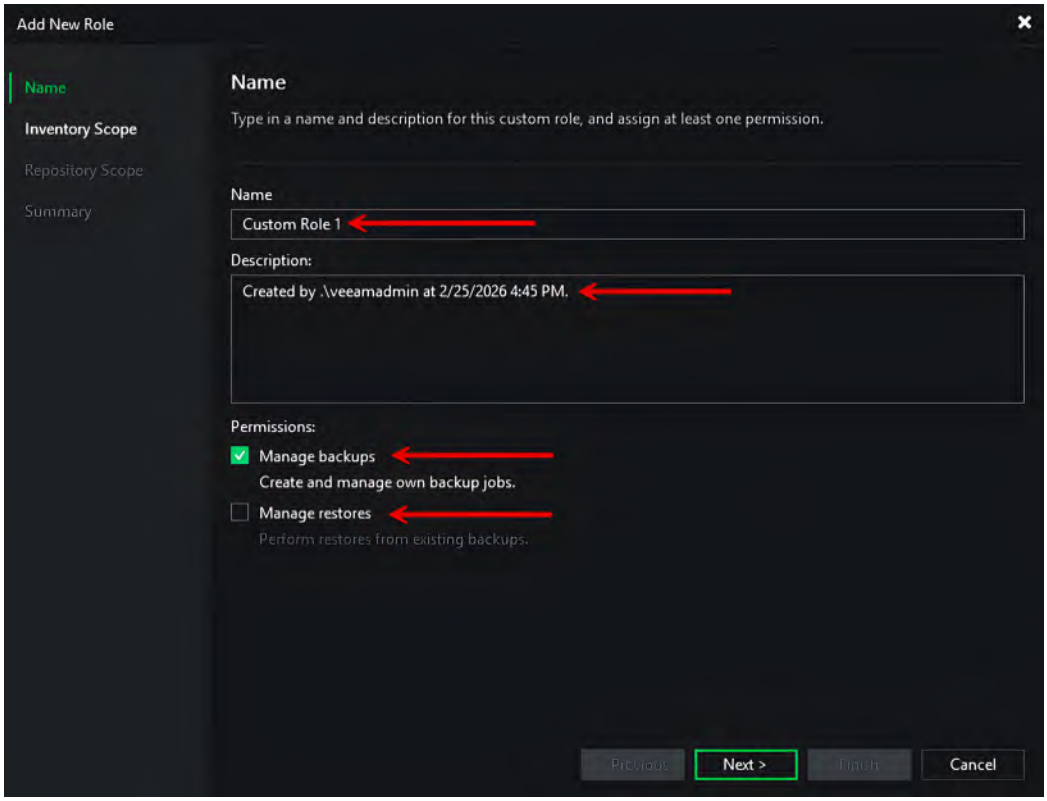


Figure 9.4: Add New Role wizard

4. In the **Inventory Scope** tab, select where the permissions will apply to, whether **Entire inventory** or **The following workloads only:** options. (In this example we will use the entire inventory option)

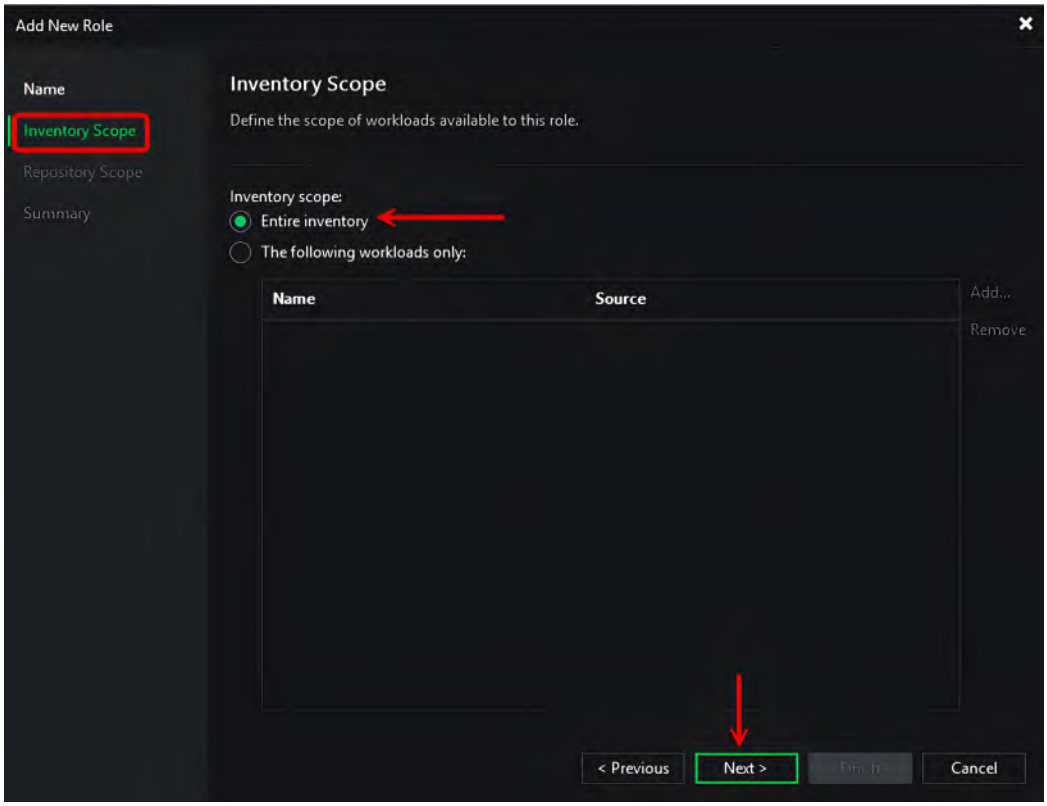


Figure 9.5: Inventory scope to assign role

5. In the **Repository Scope** tab, select where the role permissions will apply – **All backup repositories** or **The following backup repositories only:**. (In this example, we will use the **All backup repositories** option.)

Technet24

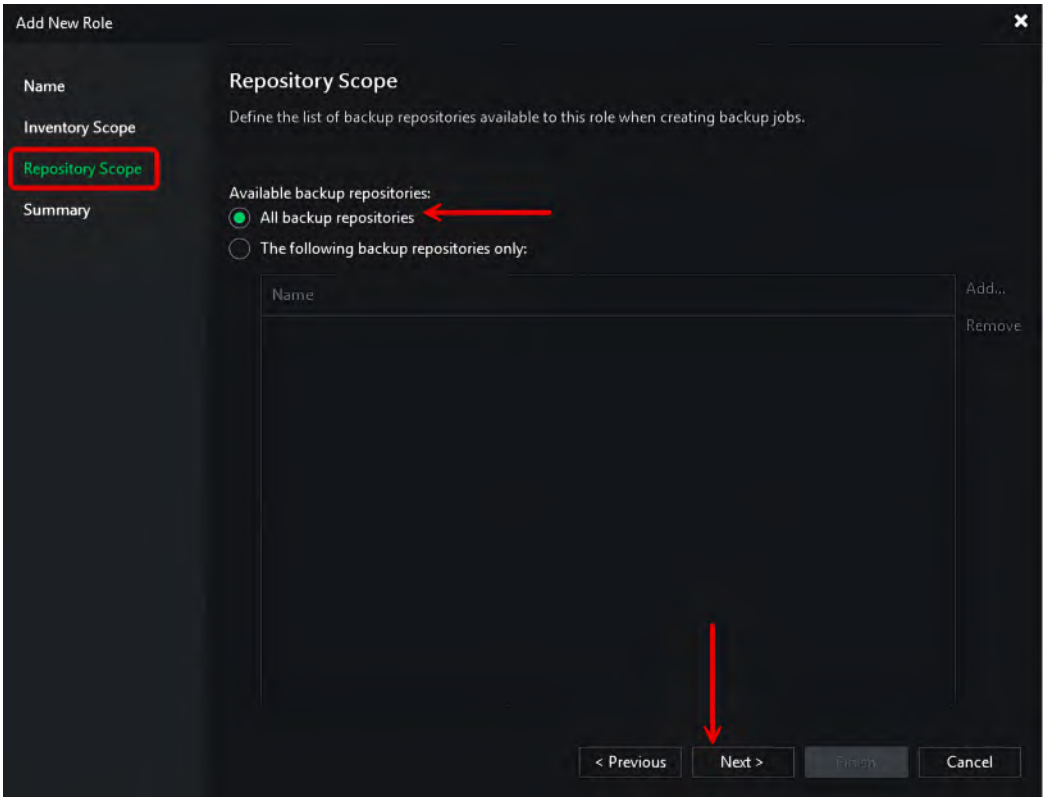


Figure 9.6: Repository scope to assign role

6. Review the summary and click **Finish** to save the custom role.

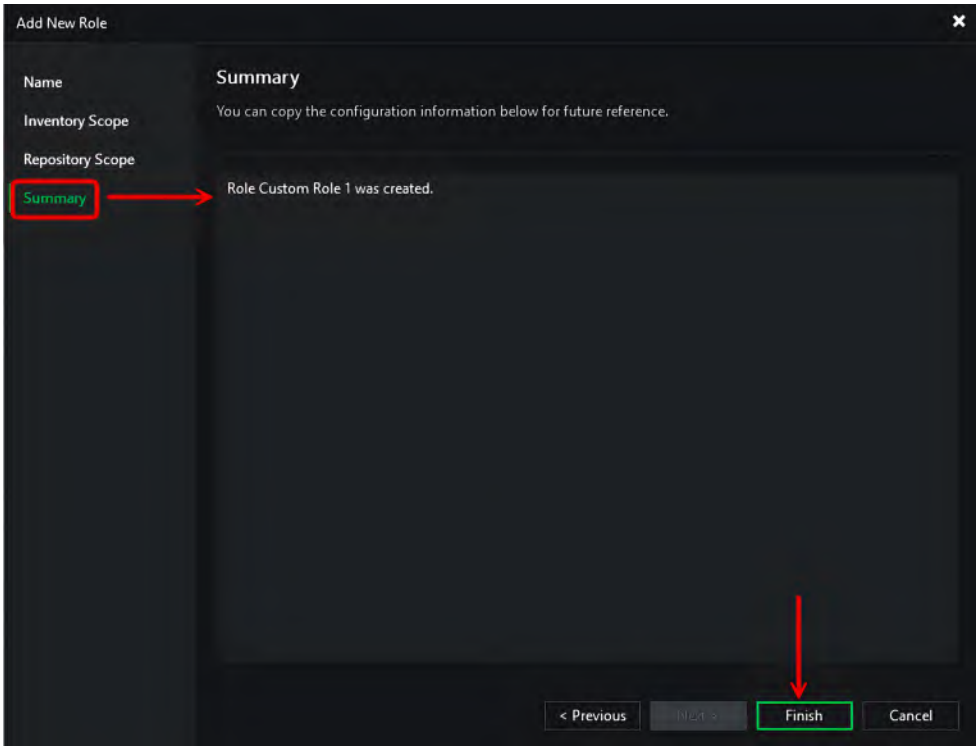


Figure 9.7: Summary of custom role

Note

You can configure highly granular permissions for both backup and restore operations across inventory and repository scopes. It is recommended to assign roles based on the principle of least privilege to avoid granting unnecessary permissions.

Permission categories

The permissions available for custom roles are organized into functional categories that map to the major operational areas of the Veeam environment. The following table summarizes the key permission categories and representative permissions within each:

Technet24

Category	Key permissions	Typical use case
Backup jobs	Create, edit, delete, start, stop, retry, and clone backup jobs	Backup administrators who manage job configurations and daily operations
Restore operations	Full VM restore, Instant Recovery, file-level recovery, application item recovery, Staged Restore	Help desk or recovery teams who perform restores on behalf of end users
Infrastructure	Add/remove managed servers, configure proxies, manage repositories, network settings	Infrastructure or storage teams responsible for the physical and virtual backup platform
Tape	Manage tape servers, media pools, vaults, and tape backup/restore jobs	Dedicated tape management teams in environments with long-term archival requirements
Configuration	General options, notification settings, traffic rules, credential management	Senior administrators who manage global backup server configuration
Monitoring & reporting	View dashboards, generate reports, access session logs, and event notifications	Management, compliance, and audit teams who need visibility without operational access

Table 9.10: Permission categories available for custom RBAC role construction

Now that we have looked at the permission categories in the preceding table, let us take a look at how we assign roles with specific scopes assigned to them.

Assigning roles and defining scope

Once custom roles have been created, they must be assigned to Active Directory users or groups. The assignment process also allows administrators to define the scope of the role, that is, which infrastructure objects the role's permissions apply to.

Role assignment process

To assign a role, follow these steps:

1. Return to the **Users and Roles** interface, switch to the **Users** tab, and click **Add**.
2. Select the Active Directory user or group, choose the role to assign (either a built-in role or a custom role), and optionally define the scope.
   The scope can be set to **All Objects** (granting the role's permissions across the entire Veeam infrastructure) or to specific objects such as individual backup jobs, repositories, or managed servers.

Scoped assignments are one of the most powerful aspects of custom RBAC. By combining a custom role with a targeted scope, you can create highly specific access patterns. For example, you might create a "*Remote Site Operator*" role with permissions to start and stop backup jobs and perform file-level restores, then assign that role to a branch office IT group with a scope limited to the backup jobs and repositories associated with their specific site.

Scope inheritance and precedence

When a user is a member of multiple Active Directory groups that have been assigned different roles, Veeam applies an additive (union) model for permissions. The effective permissions for the user are the combination of all permissions from all assigned roles. If one role grants permission and another does not, the user will have that permission. There is no explicit "deny" mechanism in the Veeam RBAC model. Access is controlled by the presence or absence of permissions in the assigned roles. An example of this would be as follows:

Imagine a user **jsmith@corp.local** is a member of two Active Directory groups, each assigned a different Veeam v13 custom role:

AD group	Veeam role assigned	Key permissions granted
GRP-Veeam-operators	Backup operator	Start/stop jobs, view job status, no restore
GRP-Veeam-restores	Restore operator	Restore VMs and files, view jobs, no job editing

Table 9.11: Example Active Directory groups for user and permissions granted

The following would be the permissions granted to the user with these two groups assigned.

Technet24

Effective permissions (union model): Because Veeam v13 applies an additive (union) model, jsmith's effective permissions are the combination of both roles.

- Start and stop backup jobs (from backup operator)
- View job status (from both roles)
- Restore VMs and files (from restore operator)
- Edit or create jobs (neither role grants this)
- Manage repositories or infrastructure (neither role grants this)

Note

Security consideration:

Because the RBAC model is additive, be careful when assigning multiple roles to the same user or group. A user who is a member of both a "*Restore Operator*" group and a "*Backup Operator*" group will effectively have the combined permissions of both roles. Review effective permissions carefully when designing your RBAC strategy, especially for users who may be members of multiple AD groups.

Real-world RBAC design patterns

The following design patterns illustrate common approaches to structuring custom RBAC roles in enterprise environments. These patterns can be adapted to suit your organization's specific requirements.

Pattern 1: Tiered administration

In this pattern, administrative responsibilities are divided into tiers based on the level of access required. Tier 1 (help desk) has restore-only access; Tier 2 (backup operations) can manage jobs and perform restores; and Tier 3 (infrastructure administration) has full access to all functions. This aligns with common ITIL-based support models.

Pattern 2: Geographic segmentation

For organizations with multiple sites, custom roles can be combined with scoped assignments to create geography-specific access. Each site's IT team receives a role that grants operational permissions (start/stop jobs, perform restores) scoped only to the backup jobs and repositories associated with their site. A central backup team retains a global administrative role for managing the overall infrastructure and cross-site operations.

Pattern 3: Separation of duties

In regulated environments, separation of duties may be required to ensure that no single individual can both create backup jobs and restore the data they backed up. Custom roles can

enforce this separation by creating distinct roles for backup operations and restore operations, assigned to different individuals or groups. This provides an audit trail that demonstrates compliance with regulatory requirements such as SOX, HIPAA, or GDPR data handling mandates.

Auditing and compliance

Custom RBAC roles integrate with the Veeam audit logging framework. All actions performed by users operating under custom roles are logged with the user's identity, the role under which the action was performed, and the specific objects affected. These audit logs can be exported for compliance reporting and integrated with SIEM solutions for centralized security monitoring.

Administrators should periodically review role assignments and effective permissions to ensure that access controls remain aligned with organizational changes. The **Users and Roles** interface provides a comprehensive view of all current assignments, and the audit log can be used to identify inactive accounts or roles that may need to be revoked.

> **Note**
>
> **Best practice - Regular access reviews:**
>
> Schedule quarterly reviews of RBAC role assignments to ensure compliance with the principle of least privilege. As personnel change roles within the organization, their Veeam RBAC assignments should be updated accordingly. Document your RBAC design decisions, including the rationale for each custom role, to facilitate audit responses and knowledge transfer.

PowerShell integration for RBAC management

For environments with many roles and assignments, PowerShell can be used to automate RBAC management tasks. The Veeam PowerShell module includes cmdlets for creating roles, assigning permissions, and managing user assignments programmatically. This is particularly useful for organizations that manage their Veeam RBAC configuration as code, allowing role definitions to be version-controlled and deployed through automated pipelines.

Key cmdlets for RBAC management include `Get-VBRUserRoleAssignment` for retrieving current role assignments, `Set-VBRUserRoleAssignment` for modifying assignments, and the various role management cmdlets for creating and configuring custom roles. These cmdlets support pipeline input, making it straightforward to perform bulk operations such as assigning a new role to all members of a specific AD group or revoking a deprecated role across all current assignees.

Technet24

The centralized updates (Veeam updater) and custom RBAC features in Veeam Backup & Replication v13.0.1 represent major improvements in operational efficiency and security governance. Centralized updates eliminate the manual, per-component patching burden and introduce HA-aware update orchestration for VSA deployments. Custom RBAC provides the granular, scope-based access control needed to align Veeam permissions with organizational roles and regulatory requirements. Together, these features enable administrators to manage larger, more complex backup environments with greater confidence and less operational overhead.

Summary

This chapter has provided you with a look into some of the new enhancements that have been made with the Veeam Backup & Replication v13 platform. We discussed the new 64-bit-only architecture and why this is now a requirement. Then, we addressed the Linux-only infrastructure and how removing Windows allows for a more secure environment. Lastly, we examined the centralized and automatic upgrade feature that is now built into the Veeam environment when using all Linux servers, and we looked at the new custom RBAC roles that you can create for assignment to users based on their role within the backup administration departments.

This chapter covered some interesting enhancements to the Veeam infrastructure to ensure you have a more secure and resilient setup. In the next chapter, we will explore more new enhancements like Postgres 17.6, the new Job Manager scalability enhancements and further refinements.

Further reading

- Veeam System Requirements - `https://helpcenter.veeam.com/docs/vbr/userguide/system_requirements.html?ver=13`
- Linux Infrastructure Components - `https://helpcenter.veeam.com/docs/vbr/userguide/linux_infrastructure.html?ver=13`
- Veeam Updates and Automatic Updates - `https://helpcenter.veeam.com/docs/vbr/em/em_update_linux.html?ver=13`
- Veeam Custom Roles - `https://helpcenter.veeam.com/docs/vbr/userguide/configure_roles.html?ver=13#custom-roles`

10

Backup Infrastructure – PostgreSQL 17.6, Improved Job Manager Scalability, and Other Changes

Veeam Backup & Replication v13 is part of the **Veeam Data Platform** and has many enhancements since the newest release of the software. This chapter will discuss **Postgres 17.6 support** and enhancements for Veeam. We will investigate the new **job manager scalability** improvements that were added. We will also discuss other enhancements including **object storage** and **immutability improvements**.

In this chapter, we're going to cover the following main topics:

- Postgres 17.6 Support
- Improved job manager scalability
- Object storage and immutability improvements

PostgreSQL 17.6 support

Veeam Backup & Replication v13.0.1 extends its configuration database compatibility to PostgreSQL 17.6, the latest maintenance release in the PostgreSQL 17 major version. This update reflects Veeam's continued commitment to supporting current, actively maintained database versions as the foundation for the VBR configuration database, which is the central repository that stores all backup infrastructure definitions, job configurations, session history, and operational metadata.

Technet24

PostgreSQL was first introduced as a supported configuration database backend in Veeam Backup & Replication v12, replacing the long-standing dependency on Microsoft SQL server express for smaller deployments and offering an open-source, license-free alternative to SQL server standard and enterprise editions in larger environments. Since that initial release, Veeam has tracked PostgreSQL's release cadence closely, with each VBR update extending support to newer PostgreSQL minor and major versions as they reach production stability.

PostgreSQL 17.6 is a maintenance release, meaning it carries no breaking changes relative to PostgreSQL 17.0. But, it does include targeted bug fixes, security patches, and reliability improvements that directly benefit long-running, write-intensive workloads like those generated by the VBR configuration database during active backup windows.

PostgreSQL 17: Key improvements relevant to VBR

Understanding what changed in the PostgreSQL 17 major release line provides you with important context for upgrading benefits VBR deployments specifically. The following improvements are most relevant to the configuration database workload profile:

- **Vacuuming and dead tuple management**: PostgreSQL 17 introduced significant enhancements to the autovacuum subsystem, including more aggressive dead tuple removal during active write periods. VBR generates substantial write activity against its configuration database during backup job execution, particularly in the BJobs, BSessions, and BStorages schema areas. This accumulates dead tuples at a rate that older PostgreSQL versions handled less efficiently when it comes to performance within the database. In PostgreSQL 17, autovacuum workers are better able to keep pace with this churn, reducing the risk of table bloat accumulating over long operational periods between manual maintenance windows.
- **Incremental sorting improvements**: Query planner improvements in PostgreSQL 17 include refined incremental sort strategies that benefit ordered result sets over large session history tables. Operators who frequently query backup session history either through the VBR console, PowerShell cmdlets, or the REST API can observe improved response times on history-intensive queries in PostgreSQL 17 compared to 15 or 16.
- **Logical replication enhancements**: For organizations that replicate the VBR configuration database to a secondary PostgreSQL instance as part of a high-availability architecture, PostgreSQL 17 introduces improved logical replication slot handling and failover slot behavior. Slots now survive planned primary switchovers more reliably, reducing the need to manually re-establish replication after intentional failover events.
- **Security and CVE remediation**: PostgreSQL 17.6 specifically addresses several CVEs present in earlier 17.x releases. Deployments in regulated environments such as

healthcare, financial services, government should treat the upgrade to 17.6 as a security patch and prioritize it accordingly rather than treating it as optional. Such organizations will have stricter security regulations and would therefore consider this a priority update.

We've discussed some of the important enhancements to PostgreSQL 17 here. Now, let us look at the supported deployment configurations that can be used with PostgreSQL 17 and Veeam Backup & Replication.

Supported deployment configurations

VBR v13.0.1 supports PostgreSQL 17.6 in the following configuration database deployment scenarios:

- **Local PostgreSQL instance (Default)**: The default installation path places PostgreSQL on the same Windows server host as the VBR server role. The VBR installer manages PostgreSQL installation, including initial cluster initialization, authentication configuration, and service registration. In this topology, PostgreSQL 17.6 is available as a selectable version during new installations and as a target version during in-place database engine upgrades.
- **Remote PostgreSQL instance**: For enterprise deployments where the configuration database is hosted on a dedicated database server, either for resource isolation, easier backup of the configuration database itself, or HA clustering, PostgreSQL 17.6 on a remote Linux or Windows host is fully supported. PostgreSQL 17.6 packages are available from the official PostgreSQL Global Development Group (PGDG) repositories for RHEL/AlmaLinux/Rocky 8–9, Ubuntu 22.04/24.04, and Debian 11/12, as well as via the EDB Windows installer.
- **Managed PostgreSQL services**: PostgreSQL 17.6 compatibility extends to managed cloud database services where the provider has made version 17 available. Consult the Veeam help center for the current matrix of tested managed service providers and their supported minor versions, as managed service version availability varies by provider and region.

That covers the supported configurations for PostgreSQL 17 with Veeam, so now let's take a look at how the installation of Veeam Backup & Replication 13 allows for a clean install of PostgreSQL 17 during the process.

Fresh installation with PostgreSQL 17.6

When performing a new installation of VBR v13.0.1, the setup wizard presents a database configuration screen that allows selection of the PostgreSQL version to be installed locally. PostgreSQL 17.6 is the recommended default for new deployments.

Technet24

The installer performs the following actions automatically when deploying a local PostgreSQL 17.6 instance:

1. Extracts the bundled PostgreSQL 17.6 Windows binaries to `%ProgramFiles%\PostgreSQL\17`.
2. Initializes a new database cluster at the configured data directory (default: `%ProgramData%\Veeam\Backup\PostgreSQL\17`).
3. Configures `pg_hba.conf` to restrict connections to localhost using `scram-sha-256` authentication.
4. Sets `postgresql.conf` parameters appropriate for the VBR workload profile, including `max_connections`, `shared_buffers`, `work_mem`, and `checkpoint_completion_target`.
5. Registers the PostgreSQL service under the `NT AUTHORITY\NetworkService` account.
6. Creates the VeeamBackup database and applies the VBR schema.

Note

No manual PostgreSQL configuration is required for local deployments. The VBR installer's defaults are tuned for the expected configuration database workload and should not be modified without guidance from Veeam support, as non-standard `postgresql.conf` values can affect VBR console responsiveness and job scheduling behavior.

For remote PostgreSQL deployments, the target PostgreSQL 17.6 instance must be prepared before running the VBR installer. Required prerequisites include the following:

- A dedicated database user with CREATEDB and CREATEROLE privileges, or an existing VeeamBackup database with the VBR service account granted ownership.
- `pg_hba.conf` entries permitting the VBR server's IP address to authenticate using scram-sha-256 or md5.
- Network connectivity on port 5432 (default) between the VBR server and the PostgreSQL host, with applicable firewall rules in place.
- SSL configured on the PostgreSQL instance if the connection traverses untrusted network segments.

Upgrading from previous PostgreSQL versions

Organizations running VBR on PostgreSQL 14, 15, or 16 can upgrade the configuration database engine to PostgreSQL 17.6 following one of two paths depending on their deployment topology.

- **Upgrade via VBR maintenance mode (local deployments)**: For local PostgreSQL deployments managed by VSA, upgrades are performed through the new automatic update mechanism, as packages downloaded for patches, hotfixes, etc. may also include PostgreSQL upgrades. If you are using the Windows installation, PostgreSQL must be upgraded manually to version 17.6.
- **Manual upgrade (remote deployments)**: For remote PostgreSQL instances, the engine upgrade must be performed by the database administrator independent of VBR. The supported upgrade path uses `pg_upgrade` with the `--link` flag where storage permits, or a dump-and-restore approach using `pg_dumpall` followed by restore into a freshly initialized PostgreSQL 17.6 cluster.
  After the remote PostgreSQL instance has been upgraded to 17.6, VBR should run and connect to the upgraded instance of Postgres automatically.
- **Upgrading from PostgreSQL 13 or earlier**: PostgreSQL 13 reached end-of-life in November 2025. Direct `pg_upgrade` from PostgreSQL 13 to 17 is supported by PostgreSQL but requires both old and new binary sets to be present on the upgrade host. If the PostgreSQL 13 binaries have been removed, a dump-and-restore migration is required. Consult Veeam KB article references in the release notes for step-by-step procedures specific to this upgrade path.

As you can see, there are a few ways in which PostgreSQL can be upgraded to the latest release. Now let's move on to the configuration database maintenance options.

Configuration database maintenance

PostgreSQL 17.6 includes autovacuum enabled by default, and the VBR installer configures autovacuum parameters appropriate for the configuration database workload. However, operators should be aware of the following maintenance considerations in long-running deployments.

Session history table growth

The BJobSessions and BTaskSessions tables are the fastest-growing tables in the VBR configuration database schema. In deployments with large numbers of jobs running on short schedules, these tables can accumulate millions of rows over months of operation. While autovacuum manages dead tuple cleanup, live row growth requires active management through the VBR session history retention settings.

Technet24

Navigate to **Main Menu | General Options | History** to configure the session retention period. Reducing retention for completed successful sessions while retaining failed session data longer is a practical approach that controls growth without sacrificing diagnostic history.

Index bloat

High-frequency insert and delete patterns on session tables can cause B-tree index bloat over time, degrading query performance on history lookups. PostgreSQL 17's improved vacuum behavior reduces the rate of index bloat accumulation, but periodic index maintenance via `REINDEX CONCURRENTLY` remains advisable for databases older than 12–18 months. This operation can be run without interrupting VBR operations when executed with the `CONCURRENTLY` option.

Connection pooling

VBR does not deploy PgBouncer or another connection pooler by default. In large-scale deployments where many concurrent backup jobs generate high configuration database connection rates, the `max_connections` parameter set during installation may require tuning. Increasing `max_connections` requires a PostgreSQL service restart; plan this during a maintenance window.

Known limitations and considerations

With every application there are some known limitations and considerations you must take into account when deploying. Let's look at some of those for PostgreSQL:

- PostgreSQL 17.6 on Windows requires Windows server 2019 or Windows server 2022. Windows server 2016 is not supported for PostgreSQL 17.x deployments.
- Mixed-version environments where a VBR v13.0.1 installation points to a PostgreSQL 16.x remote instance remain supported. Upgrading the database engine to 17.6 is recommended but not mandatory for v13.0.1 compatibility.
- VBR does not support PostgreSQL 17 beta or release candidate builds. Only the 17.6 GA release is validated for production use.
- The enterprise manager component, if deployed, shares the same configuration database and inherits the PostgreSQL version requirements of the VBR server. Ensure enterprise manager is at v13.0.1 before upgrading the shared configuration database engine.

You must check these prerequisites and operating system compatibility. Now let's look at another improvement with the job manager in Veeam Backup & Replication 13.

Improved job manager scalability

Veeam Backup & Replication v13.0.1 introduces a set of targeted improvements to the **job manager** subsystem aimed at organizations operating at scale where environments are characterized by hundreds to thousands of defined jobs, high rates of concurrent job execution, dense proxy and repository infrastructure, and backup windows that demand precise scheduling fidelity across large job populations.

Job manager is the core scheduling and orchestration engine within the Veeam backup service. It is responsible for evaluating job schedules, queuing jobs awaiting resource availability, dispatching jobs to available proxies and repositories, tracking execution state across all active and recently completed sessions, and enforcing the resource allocation policies defined by the backup administrator. In smaller deployments, job manager operates well within comfortable headroom. In large enterprise and service provider environments, however, the cumulative demands of scheduling, dispatching, and tracking thousands of concurrent and sequential jobs can introduce latency, scheduling drift, and resource contention that degrades the reliability of backup windows.

The improvements in v13.0.1 address three distinct layers of the job manager: the internal scheduling engine that evaluates and queues jobs, the dispatch layer that allocates proxy and repository resources to queued jobs, and the state tracking system that maintains session metadata for active and recently completed jobs. Together, these changes extend the practical upper bound of job population and concurrency that a single VBR server can manage reliably.

Job manager architecture

To understand the nature of the scalability improvements in v13.0.1, it is useful to familiarize with the operational model of the job manager and where stress accumulates in large deployments.

The Veeam backup service maintains an in-memory representation of all defined jobs, their schedule definitions, their current state, and the resource constraints that govern when they can be dispatched. At each scheduling interval, the service evaluates which scheduled jobs are due to start, checks whether their required resources such as proxies, repositories, and WAN accelerators have available capacity within the defined concurrency limits, and either dispatches the job immediately or places it in the wait queue.

In large environments, this evaluation loop, sometimes called the **scheduling tick**, must process a substantial population of job definitions at each interval. If the evaluation loop cannot complete within the scheduling interval, scheduling drift begins to occur. Jobs that were due to start at a given time instead start at a later interval, with the delay compounding as job population grows. Simultaneously, the dispatch layer must track all active job sessions

Technet24

and update resource utilization counters in near real-time to ensure that resource limits are accurately enforced.

Prior to v13.0.1, the scheduling tick and resource tracking operations shared a single execution context within the Veeam backup service process. This design was appropriate for environments up to several hundred jobs but introduced measurable latency in environments with job populations above approximately 1,500 defined jobs, particularly when a large proportion of those jobs were concurrently active or in the wait queue.

Scheduling engine improvements

Job scheduling within Veeam Backup & Replication that ensures jobs run successfully. Let's explore the enhancements and improvements in this area.

Parallel schedule evaluation

The most significant internal change in v13.0.1 is the introduction of **parallel schedule evaluation** within the scheduling tick. In prior versions, job schedule conditions including time-based triggers, after-job chaining conditions, and backup window restrictions were evaluated sequentially across the full job population. In v13.0.1, the evaluation of independent jobs within each scheduling tick is distributed across a thread pool, with the pool size automatically scaled to the number of logical processors available on the VBR server host.

The practical effect is that the scheduling tick duration scales favorably with both job population size and available CPU resources. Benchmark testing conducted by Veeam engineering on reference hardware (a VBR server with 16 logical cores and 2,000 defined jobs) showed scheduling tick durations reduced by approximately 60–70% compared to v12.x under equivalent workloads. This directly reduces scheduling drift in high-density environments and allows the scheduling interval to remain tight even as job populations grow.

Schedule condition caching

Time-based and window-based schedule conditions are now evaluated against a cached representation of the current time window rather than re-evaluated independently for each job. In deployments with large numbers of jobs sharing common backup windows, this eliminates redundant computation and reduces the CPU overhead of the scheduling tick during periods of high job density. This will typically be at the start and end of primary backup windows when large populations of jobs transition simultaneously between eligible and ineligible states.

Improved after-job chaining

After-job chaining, where a job is configured to start upon completion of a predecessor job, previously required the scheduling engine to re-evaluate the full dependency graph on each tick to detect newly satisfied chain conditions. In v13.0.1, chain condition evaluation is event-

driven: when a job completes, its completion event directly triggers evaluation of any dependent jobs rather than waiting for the next scheduling tick. In workflows with deep job chains or large fan-out topologies, this eliminates latency between steps in the chain that previously accumulated across multiple scheduling intervals.

Dispatch layer improvements

As with any software, there is a layer that manages everything and sends tasks to different components, Veeam is no different here from the way tasks are performed and sent to proxy servers for example. We'll now look at some improvements in this area.

Resource utilization tracking decoupled from scheduling

Prior to v13.0.1, resource utilization counters which tracked how many concurrent tasks each proxy and repository was currently serving, were updated synchronously within the scheduling tick. In large environments with many active jobs, this meant the scheduling tick was blocked waiting for utilization counter updates to complete before it could evaluate the next batch of jobs.

In v13.0.1, resource utilization tracking is decoupled into a dedicated background process that updates counters asynchronously. The scheduling tick reads the most recently committed utilization snapshot rather than triggering live updates. The snapshot refresh rate is configurable and defaults to a sub-second interval, maintaining effective resource limit enforcement while eliminating the scheduling tick's dependency on synchronous utilization queries.

Proxy selection algorithm optimization

When multiple proxies are eligible to service a job, the proxy selection algorithm must evaluate each eligible proxy's current load, transport mode compatibility, network proximity to the source datastore, and any explicit proxy assignment rules. In v13.0.1, the proxy selection algorithm has been reworked to short-circuit evaluation as soon as a suitable proxy is identified in common scenarios, rather than exhaustively scoring all eligible proxies before selecting the highest-scoring candidate. For deployments with large proxy pools (twenty or more proxies,) this reduces dispatch latency for straightforward job types while preserving full multi-criteria scoring for complex proxy selection scenarios involving transport mode preferences and explicit affinity rules.

Repository queue management

Repository throughput limits are enforced by a queue at the repository level. Jobs that would exceed the repository's configured maximum concurrent tasks are held in the wait queue until capacity becomes available. In prior versions, the repository queue was evaluated as part of the main scheduling tick.

Technet24

In v13.0.1, the repository queue has its own dedicated evaluation loop, allowing jobs waiting for repository capacity to be dispatched immediately when a slot becomes available without waiting for the next main scheduling tick. In large environments where repository contention is the primary cause of job queueing, this improvement directly reduces the time jobs spend in the wait state.

Session state management

Veeam manages each backup session and there have been significant improvements made within this area. These are covered in the following sub-sections.

Reduced memory footprint per active session

Each active job session maintains an in-memory state object within the Veeam backup service that tracks task progress, resource assignments, log messages, and intermediate results. In prior versions, these state objects were allocated at the maximum possible size to accommodate the largest job configurations. In v13.0.1, session state objects are allocated dynamically based on the actual job configuration, with memory reclaimed progressively as tasks within the session complete.

For large environments with many concurrent sessions, this reduces the working set memory consumption of the Veeam backup service proportionally to the average number of VMs or workloads per job. In environments with many small jobs such as a common service provider pattern where each protected tenant has one or a small number of jobs, the per-session memory reduction is most pronounced. Veeam engineering measurements indicate a 20–35% reduction in Veeam backup service working set memory under sustained high-concurrency load compared to v12.x.

Session history write optimization

As job sessions are complete, their full session records including per-task results, log entries, and duration metadata are written to the configuration database. In prior versions, this write occurred synchronously on the session completion path, adding latency to session teardown that accumulated when many sessions completed within a short time window (typically at the end of a backup window as jobs finish in rapid succession).

In v13.0.1, session history writes are batched and dispatched asynchronously via a dedicated write queue. Sessions are marked as completed in the in-memory state immediately upon completion, allowing resources to be released and dependent jobs to be triggered without waiting for the database write to be complete. The database write follows within a configurable grace period (default: 5 seconds). This does not affect data integrity. The session record is committed before the in-memory state object is released, but it eliminates the write latency from the critical path of session teardown and resource release.

Large-scale deployment guidance

Veeam deployments within a large infrastructure environment require some consideration prior to installation. We will now examine some of the guiding principles to follow.

VBR server sizing

The scalability improvements in v13.0.1 are designed to extend the viable job population and concurrency ceiling for a single VBR server instance. However, the benefits are only realized if the VBR server host is sized appropriately. The following minimum specifications are recommended for environments with job populations above 1,000 defined jobs:

- **CPU**: 16 or more logical cores. The parallel scheduling engine scales with available cores; under-provisioned CPU is the most common constraint in high-density deployments.
- **RAM**: 32 GB or more. Session state memory improvements reduce per-session overhead, but the aggregate working set of a large active job population is substantial.
- **Storage (OS/application disk)**: SSD-backed storage for the VBR application binaries and configuration database (if local). Scheduling tick performance is sensitive to configuration database query latency, rotational disk increases tick duration under load.
- **Network**: A dedicated NIC or NIC team for backup traffic is strongly recommended in environments where the VBR server also participates in data movement (e.g., hot-add transport mode).

Backup window design

The scheduling improvements in v13.0.1 reduce drift, but backup window design remains the primary lever for managing scheduling fidelity at scale. The following are some best practices for large environments:

- Distributing job start times across the backup window using offset start times or staggered schedules rather than starting large job populations simultaneously.
- Using backup window restrictions (permitted hours) to define outer bounds rather than relying on precise start times for SLA enforcement.
- Leveraging after-job chaining for workflows where sequencing matters, taking advantage of the event-driven chain triggering introduced in v13.0.1.
- Avoiding deeply nested job chains (more than 5 levels) in very large environments, as chain evaluation overhead scales with chain depth regardless of the evaluation model improvements.

Technet24

Now that we've covered some of the considerations you need to look at when deploying at a large scale, let's discuss the ability to monitor the health of your environment next.

Monitoring scalability health

VBR v13.0.1 adds new metrics to the built-in statistics views relevant to job manager scalability. These metrics are accessible via the VBR REST API and the Veeam ONE integration:

- **Scheduling tick duration**: The time taken to complete each evaluation cycle of the scheduling engine. A healthy value is below 500 ms; sustained values above 1,000 ms indicate the scheduling engine is under pressure and scheduling drift may be occurring.
- **Wait queue depth**: The number of jobs currently in the wait state, segmented by wait reason (awaiting proxy capacity, awaiting repository capacity, awaiting global concurrency slot, within backup window restriction). Sustained high queue depth for a specific reason identifies the bottleneck limiting throughput.
- **Session teardown latency**: The time between a session completing and its resources being released back to the available pool. Values consistently above 10 seconds in the post-v13.0.1 baseline may indicate database write contention affecting the asynchronous session history write queue.

These metrics can be surfaced in Veeam ONE dashboards and used to establish baseline values for a given environment, enabling proactive identification of scalability constraints before they manifest as visible backup window failures.

Known limitations and considerations

As with all large environments, there are always the following limitations and considerations to take into account:

- The parallel scheduling engine benefits scale with CPU core count up to 32 logical cores. Environments with VBR server hosts exceeding 32 cores will not see proportional additional benefit from further core increases within the scheduling subsystem specifically.
- Asynchronous session history writes introduce a brief window (default: 5 seconds) during which a completed session's record may not yet be visible in the VBR console's session history view. This is cosmetic and does not affect job execution or resource release.
- The increased default global concurrent job limit of 200 applies to new installations only. Upgraded environments must manually adjust this value if higher concurrency is desired.

- In environments using Veeam ONE for job scheduling monitoring, the Veeam ONE collector must be at v13.0.1 or later to consume the new scheduling tick duration and wait queue depth metrics introduced in this release.
- Enterprise manager deployments that aggregate multiple VBR servers do not aggregate scheduling metrics across servers; each VBR server's scheduling health must be monitored independently.

Let's move on to the improvements made to the object storage and immutability.

Object storage and immutability improvements

Veeam Backup & Replication v13.0.1 delivers a broad set of enhancements across two closely related capability areas: **object storage integration** and **backup data immutability**. While these areas address distinct technical concerns, they are architecturally coupled in modern backup environments. Immutability is increasingly delivered through object storage primitives such as S3 Object Lock and Azure Blob immutability policies, and the reliability of immutability guarantees depends directly on the correctness and consistency of the object storage operations that create and manage protected data.

The improvements in v13.0.1 span the full lifecycle of object storage usage within VBR: initial data movement from performance tier to capacity tier, direct backup-to-object workflows, archive tier offload, metadata consistency during failure recovery, and the management of immutability lock periods across heterogeneous storage providers. For organizations that have adopted object storage as a primary or secondary backup target and particularly those using immutability as a ransomware recovery safeguard, this release addresses several operational pain points identified in prior versions and extends immutability support to additional storage providers and deployment scenarios.

Object storage enhancements

There have been many enhancements and improvements with object storage, and we will discuss them now.

Capacity tier offload performance

The capacity tier in a scale-out backup repository (SOBR) allows backup data to be moved from performance tier extents to object storage after a configurable period, reducing the long-term cost of retaining backup data while keeping recent restore points on fast local storage. The offload process in prior versions was constrained by a single-threaded upload pipeline per SOBR. Each offload job serialized the upload of backup files to object storage, limiting throughput to the upload speed achievable by a single stream regardless of available network bandwidth.

Technet24

In v13.0.1, the capacity tier offload pipeline is parallelized at the file level. Multiple backup files are uploaded concurrently within a single offload operation, with the degree of parallelism configurable per SOBR (default: 4 concurrent uploads). On high-bandwidth connections to object storage such as direct peering, ExpressRoute, or Direct Connect, this change alone can reduce offload window duration by 50–75% for repositories with large populations of files eligible for offload.

The parallelization applies equally to data retrieval. When restore operations require data to be recalled from the capacity tier, multiple files are downloaded concurrently, reducing the time-to-restore for jobs whose restore points span multiple object storage files.

Multipart upload reliability

Large backup files uploaded to object storage are split into parts and uploaded using the provider's multipart upload API, which allows interrupted uploads to be resumed rather than restarted from the beginning. In prior versions, VBR did not persist in the state of in-progress multipart uploads across service restarts. If the Veeam backup service was restarted during a capacity tier offload operation, any in-progress multipart uploads were abandoned and restarted from the beginning on the next offload cycle, with the incomplete multipart upload artifacts remaining in the bucket until manual cleanup.

In v13.0.1, multipart upload state persists to the configuration database at each completed part boundary. If the Veeam backup service restarts during an offload, the upload resumes from the last committed part boundary rather than restarting. This prevents the accumulation of incomplete multipart upload artifacts and reduces the time required to complete offload operations that span service maintenance windows.

A companion improvement introduces automatic cleanup of orphaned multipart upload artifacts, incomplete uploads that predate the v13.0.1 upgrade or resulted from abnormal terminations in prior versions. During the first capacity tier maintenance cycle after upgrading to v13.0.1, VBR inventories the target bucket for multipart upload artifacts older than a configurable threshold (default: 48 hours) and aborts them, recovering the associated storage capacity.

Archive tier improvements

VBR's archive tier extends the SOBR tiering model to support deep archive storage classes such as Amazon S3 Glacier Instant Retrieval, S3 Glacier Flexible Retrieval, S3 Glacier Deep Archive, Azure Archive, and equivalent classes from other providers for long-term retention of backup data at minimum cost. Archive tier improvements in v13.0.1 address two significant operational limitations present in prior releases.

Selective archive tiering by retention policy

Prior versions moved all capacity tier data eligible for archive based solely on age, without the ability to differentiate between grandfather-father-son (GFS) restore points that must be retained for compliance purposes and standard recovery restore points that have aged out of the active retention window. In v13.0.1, archive tier policies can be configured to target GFS restore points specifically, allowing weekly, monthly, and yearly GFS points to be promoted to archive storage independently of standard daily restore points. This enables organizations to define cost-optimized retention architectures where compliance-tier data is archived at deep storage pricing while recent operational restore points remain on capacity tier object storage for faster retrieval.

Archive retrieval cost estimation

Restoring from archive storage classes incurs retrieval fees that vary by provider, storage class, and retrieval speed tier. Prior versions of VBR initiated archive retrievals without presenting a cost estimate, creating risk of unplanned retrieval charges in production restore scenarios. In v13.0.1, restore operations that require archive tier data retrieval display an estimated retrieval cost and duration based on current provider pricing data and the selected retrieval speed. And they also require explicit confirmation before initiating the retrieval request. The cost estimation data is updated periodically by the VBR service and can be refreshed manually from the repository settings if pricing has changed since the last update.

Direct backup to object storage

Direct backup to object storage where VBR writes backup data directly to an S3-compatible bucket or Azure Blob container without an intermediate performance tier extent, was introduced in VBR v12 as a primary backup target option for environments that prioritize cost and simplicity over restore performance. In v13.0.1, several limitations of the initial implementation are addressed.

Synthetic full backup support

Prior versions did not support synthetic full backup creation for direct-to-object backup jobs. Administrators were required to schedule periodic active full backups, which consumed source-side resources and transmission bandwidth equivalent to a full backup even when incremental data volume was small. In v13.0.1, synthetic full backup is supported for direct-to-object targets on providers that support server-side copy operations (S3 CopyObject, Azure Blob Copy). The synthetic full is assembled by issuing server-side copy instructions for the relevant data blocks, without transferring data through the VBR proxy. This reduces the resource impact of periodic full backup creation in direct-to-object configurations significantly.

Technet24

Immutable backup chain integrity verification

When immutability is enabled on a direct-to-object backup repository, VBR v13.0.1 introduces a scheduled integrity verification job that confirms all files in the backup chain carry active, unexpired object lock holds. The verification job runs independently of the backup schedule and reports any restore points where the object lock has expired or was not applied correctly, allowing administrators to identify and remediate gaps in immutability coverage before they become relevant in a recovery scenario.

Immutability enhancements

With all the object storage enhancements, there have been some significant enhancements when it comes to immutability. The following sub-sections will discuss the same.

S3 object lock improvements

S3 object lock is the primary immutability mechanism for object storage targets in VBR. It prevents backup data from being deleted or overwritten for a defined retention period, providing a ransomware-resistant recovery layer that survives even compromised VBR server credentials.

S3 object lock operates in two modes: **governance mode**, where users with the `s3:BypassGovernanceRetention` IAM permission can remove locks, and **compliance mode**, where locks cannot be removed by any user including the bucket owner, until the retention period expires. Prior versions of VBR applied object lock without clearly distinguishing mode selection in the UI, leading to environments where operators believed they had compliance-mode protection when governance mode had been configured.

In v13.0.1, the object lock mode is a required, explicit selection during repository configuration with a clear description of the operational and security implications of each mode. Existing repositories configured with object lock display a banner identifying their current mode and, where the current credentials have sufficient permissions, offer a one-way migration path from Governance to Compliance mode.

Lock period alignment with retention

In prior versions, the object lock retention period applied to each uploaded backup file was set to a fixed administrator-defined value at upload time, independent of the backup job's retention policy. This created scenarios where the object lock period expired before the backup job's retention policy would have removed the restore point, leaving a window where data could be deleted before it aged out of the retention window. Or conversely, where the lock period outlasted the retention policy, preventing storage reclamation after restore points had been intentionally retired.

In v13.0.1, the object lock retention period is calculated dynamically at upload time based on the backup job's configured retention policy, including GFS retention tiers. Daily restore points receive lock periods aligned to the daily retention setting. GFS weekly, monthly, and yearly restore points receive lock periods aligned to their respective retention horizons. As GFS restore points are promoted through retention tiers, the lock period is extended automatically using the object lock `ExtendObjectRetention` API call.

Immutability status in the VBR console

The VBR console now surfaces object lock status at the restore point level. In the **Backups | Object Storage** view, each restore point displays its lock expiration date and mode (governance or compliance). Restore points with expired or missing locks are flagged with a warning indicator, allowing administrators to identify immutability gaps without querying the storage provider directly.

Azure Blob immutability improvements

Azure Blob storage immutability policies operate differently from S3 object lock. They are applied at the container level rather than the object level, and they can be configured in either time-based retention policy mode or legal hold mode. VBR v13.0.1 extends its Azure Blob immutability support in two areas: **container-level policy validation** and **immutability policy renewal**.

Container-level policy validation

Prior versions did not validate the Azure Blob immutability policy configuration during repository setup or during backup operations. If the target container's immutability policy had expired or had not been configured correctly, VBR would write backup data to the container without immutability protection and report success, with no indication that the immutability guarantee was absent. In v13.0.1, VBR validates the container's immutability policy state at the start of each backup job and at the start of each capacity tier offload operation. Jobs targeting containers with invalid, expired, or absent immutability policies generate a warning in the job session log and optionally raise an alert in Veeam ONE if the integration is configured.

Immutability policy renewal

Azure time-based retention policies have a finite duration and must be extended before they expire to maintain continuous immutability coverage. Prior versions provided no tooling for immutability policy renewal within VBR. Administrators were required to manage policy extension manually through the Azure portal or Azure CLI. In v13.0.1, VBR monitors the expiration dates of immutability policies on managed Azure Blob repositories and raises alerts when a policy is within a configurable lead time of expiration (default: 30 days). For repositories where the configured service principal has the storage blob data owner role, VBR

Technet24

can optionally extend the policy automatically, aligning the new expiration date to the current retention policy of the associated backup job.

Hardened Linux repository immutability

The hardened Linux repository provides immutability for backup data stored on local Linux-based storage using `chattr +i filesystem-level immutability` flags. It is the primary immutability mechanism for on-premises backup data that does not use object storage. VBR v13.0.1 introduces two improvements to hardened repository operation: **immutability period per backup job** and **immutability verification reports**.

Immutability period per backup job

Prior versions applied a single immutability period to all backup data written to a hardened repository, configured at the repository level. This prevented administrators from applying different immutability durations to jobs with different retention requirements sharing the same hardened repository. In v13.0.1, the immutability period can be configured at the backup job level for jobs targeting a hardened repository, overriding the repository-level default. The repository-level setting remains as a floor: job-level immutability periods cannot be set below the repository minimum, preventing inadvertent reduction of immutability coverage for sensitive workloads.

Immutability verification reports

Veeam ONE v13.0.1 introduces a scheduled immutability verification task for hardened repositories that confirms all expected backup files carry active immutability flags. The verification task runs on a configurable schedule (default: daily) and produces a report accessible from the repository properties view. Files where immutability flags are absent or where the flag application timestamp is inconsistent with the backup chain timeline are reported as anomalies requiring investigation. This provides an auditable verification trail suitable for regulatory compliance reporting in environments where immutability is a control requirement.

Scale-out backup repository operational changes

The capacity tier improvements in v13.0.1 require a one-time metadata migration for existing SOBRs that have previously offloaded data to object storage. This migration is performed automatically during the first maintenance window after upgrading to v13.0.1 and updates the internal metadata structures used to track multipart upload state and per-file lock periods. The migration is non-disruptive, that is, backup and restore operations against the SOBR continue normally during the migration. The duration of the migration scales with the number of objects in the capacity tier. Environments with large object storage populations should anticipate migration durations of up to several hours.

During the migration period, capacity tier offload operations are paused. Restore operations that require capacity tier data continue to function normally. Offload operations resume automatically upon migration completion.

Administrators should verify that adequate free space exists on performance tier extents to absorb data that would normally be offloaded during the migration window. In environments where the offload schedule is aggressive and performance tier extents operate near capacity, scheduling the v13.0.1 upgrade during a period of lower backup activity is advisable.

Known limitations and considerations

As with any object storage and immutability, take the following limitations and considerations into account:

- Parallel capacity tier offload (multiple concurrent uploads) requires the object storage provider to support concurrent multipart upload sessions to the same bucket. All major S3-compatible providers support this. Some on-premises S3-compatible appliances may have per-bucket concurrency limits that constrain the benefit. Verify provider concurrency limits before increasing the parallel upload setting above the default.
- Synthetic full backup for direct-to-object targets using server-side copy requires the VBR service account to have the `s3:CopyObject` (or equivalent) permission on the target bucket. Environments using restrictive bucket policies should verify this permission is present before enabling synthetic full schedules.
- Lock period dynamic alignment with GFS retention requires the backup job's GFS retention settings to be configured in VBR. Jobs using only daily retention without GFS tiers will have lock periods aligned solely to the daily retention setting.
- The immutability verification report for hardened repositories requires SSH connectivity from the VBR server to the hardened repository host at the time the verification task runs. Environments with strict SSH access controls should ensure the VBR service account's SSH key remains authorized on the hardened repository host.
- Azure Blob immutability policy auto-renewal requires the storage blob data owner role on the target storage account. Environments where VBR operates with a less-privileged service principal must manage policy renewal manually or via external automation.
- The multipart upload artifact cleanup scan performed after upgrading to v13.0.1 may generate a temporary increase in API call volume against the target bucket. Environments with object storage providers that enforce API rate limits should monitor for throttling during the first maintenance window post-upgrade.

Technet24

Summary

This chapter has provided you with an overview of some great enhancements to the Veeam Backup & Replication v13 platform. We discussed the new PostgreSQL 17.6 support and enhancements, including upgrade paths. We also discussed the improved job manager scalability enhancements and other enhancements including both object storage improvements and Immutability improvements.

This chapter has shown you some of the great enhancements that come with Veeam Backup & Replication v13. In the next chapter, we will cover the new Veeam ONE v13 and what is new within this release.

Further reading

PostgreSQL Server Upgrade Steps — `https://www.veeam.com/kb4386`

11

Veeam ONE v13

Veeam ONE v13 is part of the **Veeam Data Platform**, the newest release of Veeam which allows you to monitor all your workloads, including **Veeam Backup & Replication**, **Veeam Backup for Microsoft 365**, and your **Virtual Infrastructure**. This chapter will discuss what is new with monitoring in Veeam ONE v13. We will explore the architectural changes with Veeam ONE v13 and what is now required. Lastly, we will look at the new dashboard, **Veeam intelligence** and reporting changes.

In this chapter, we're going to cover the following main topics:

- What's new with Veeam ONE v13
- Veeam ONE v13 - Architecture changes
- New dashboard, Veeam intelligence, and reporting changes

Technical requirements

To follow along with this chapter, please install Veeam ONE v13. You will require the following:

- You must have deployed a **Windows 2016/2019/2022/2025** server with the necessary disk space to install the application (**2012/2008 R2 SP1** is no longer supported). **Windows 10** and other modern **Windows** desktop operating systems are also supported. Please refer to the following website for the Operating System requirements: `https://helpcenter.veeam.com/docs/one/userguide/system_requirements.html?ver=13`

Technet24

- You must have downloaded the latest **ISO file** from `www.veeam.com`, which requires registration on the site and allows you to obtain a trial license. At the time of this writing, version `13.0.1.6168` is the current release.
- *Veeam Documentation* website: `https://helpcenter.veeam.com/docs/one/userguide/about_one.html?ver=13`

What's new with Veeam ONE v13

With the release of Veeam ONE v13, there are many product enhancements and changes, including some architectural changes. In this section, we will examine the same.

Threat center – The new default dashboard

The most immediate visible change in Veeam ONE v13 is the new threat center dashboard, which replaces the traditional infrastructure summary as the first screen displayed when logging into the Veeam ONE web console. This reflects a deliberate design shift, whereas earlier versions of Veeam ONE presented backup job status and infrastructure health as the primary concern, v13 leads with security posture and cyber-resilience.

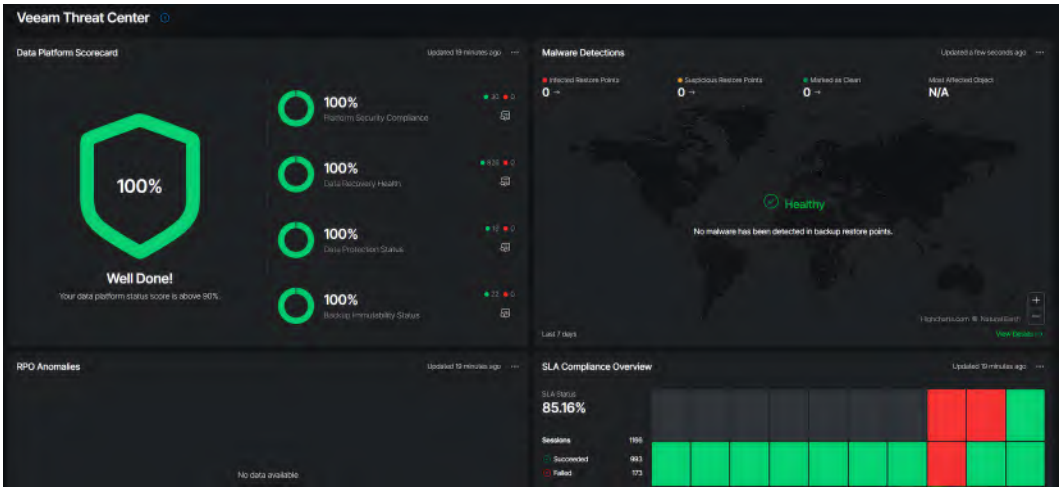


Figure 11.1: Veeam Threat center dashboard in Veeam ONE WebUI

The threat center consolidates four critical areas into a single screen, giving administrators an at-a-glance summary of where the backup environment stands regarding security, recovery readiness, and malware exposure.

Data platform scorecard

The central widget on the threat center is the **Data Platform Scorecard**, which aggregates four independent health metrics into a unified score. Each metric is displayed as a percentage, making it easy to identify which area requires attention without drilling into individual reports.

Scorecard metric	What it measures
Platform security compliance	Checks that backup infrastructure components meet security configuration baselines
Data recovery health	Validates that recent restore point tests (SureBackup) have completed successfully
Data protection status	Measures the percentage of protected workloads that have met their defined RPO
Backup immutability status	Reports on the proportion of backup data stored with immutability enabled

Table 11.1: Data platform scorecard metrics

When Coveware Recon integration is enabled, detected malware incidents are also surfaced in the scorecard, categorized by severity and suspicion level, providing a single entry point for ransomware triage.

Malware detections widget

The **Malware Detections** widget displays confirmed and suspected malware threats in the environment, visualized by geographic location and categorized by infection severity. Administrators can configure the time range for displayed detections and quickly navigate to the most affected objects. The widget also highlights clean restore points, giving recovery teams a fast path to identifying which backups can be used safely for restoration.

RPO anomalies widget

The **RPO Anomalies** widget identifies workloads that have missed their defined recovery point objective, meaning that they have not been backed up within the expected timeframe. This widget is highly configurable:

- Scope can be filtered to specific workload types: VMs, physical computers, unstructured data, cloud instances, or enterprise applications.

Technet24

- RPO thresholds are customizable per workload type.
- Wildcard-based exclusion rules allow known exceptions to be filtered out, reducing noise.

SLA compliance overview

The **SLA Compliance Overview** provides a heatmap of backup job success percentages over a configurable historical period. Cells in the heatmap transition from green through yellow to red as success rates fall below configurable SLA thresholds. This gives operations teams an immediate visual indication of which backup jobs have been consistently unreliable, without needing to page through job logs manually.

Redesigned dashboard catalog

The **Dashboard Catalog**, the library where administrators browse and launch dashboards, has been completely redesigned. The previous large-tile layout has been replaced with a more structured, navigable design that supports three UI themes, available in both light and dark modes. This is the first time Veeam ONE has offered a dark mode option.

Dashboard catalog offers the following features:

- Administrators can star frequently used dashboards to pin them at the top of the catalog.
- Three theme options support personal preference and display environment requirements (e.g., NOC screens benefit from dark mode).
- A global search from the login screen spans dashboard names, report titles, and widget descriptions.

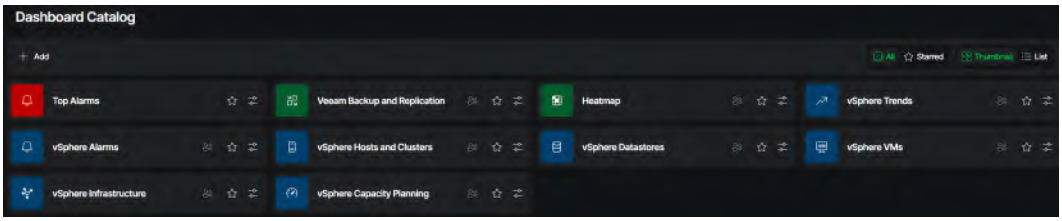


Figure 11.2: New dashboard catalog page

New and enhanced alarms

Veeam ONE v13 introduces several new alarm categories alongside enhancements to the alarm engine's integration capabilities. The following alarms are new in v13 and v13.0.1.

Application backup policy state alarm (v13.0.1)

Prior to v13.0.1, standalone Veeam plug-ins for enterprise applications were invisible to Veeam ONE and there was no way to alert on their backup policy state. The new **application backup policy state** alarm closes this gap. It tracks protection status for the following standalone plug-ins:

- Veeam Plug-in for Microsoft SQL Server
- Veeam Plug-in for SAP HANA
- Veeam Plug-in for Oracle RMAN
- Veeam Plug-in for SAP on Oracle
- Veeam Plug-in for IBM Db2
- Veeam Plug-in for SAP MaxDB

VBR high availability health alarms (v13.0.1)

For organizations deploying Veeam Backup & Replication in a High Availability (HA) configuration using VSA, new alarms provide continuous health monitoring of the active/passive cluster. These alarms can detect health degradations and failover conditions and include an automated remediation action that can trigger a failover when Veeam ONE determines the primary node is unhealthy. This enables policy-driven failovers for both planned maintenance and unplanned outages.

Coveware recon integration alarms (v13.0.1)

When Coveware Recon Scanner 3.0 is integrated with Veeam ONE, each incident flagged by Recon generates an actionable alarm in Veeam ONE. Recon monitors adversary behavior including brute force login attempts, suspicious file activity, and unexpected network connections. Filtering options reduce alarm fatigue by suppressing known benign patterns. Alarms can be routed to ServiceNow, SNMP traps, Syslog, email, or custom scripts.

VBR authentication anomaly alarms (v13.0.1 Patch 1)

New alarms capture both successful and failed sign-in activity on Veeam Backup & Replication servers. Veeam ONE surfaces trends and spikes in authentication events and correlates them with other backup events to identify risky patterns like repeated authentication failures followed by privilege escalation or suspicious job launches. These alarms integrate with the same notification channels as other Veeam ONE alarms (ServiceNow, SNMP, Syslog, email, custom scripts) and are designed to support audit and regulatory compliance requirements.

New reporting engine

One of the most significant infrastructure changes in Veeam ONE v13 is the complete replacement of the SQL server reporting services (SSRS) engine with a modern JavaScript/

Technet24

HTML5-based reporting framework. SSRS was a Microsoft technology with a fixed rendering pipeline that limited interactivity and placed heavy load on the SQL server database during report generation. The new engine eliminates these constraints.

Key capabilities of the new engine

Table 11.2 consolidates the main capabilities of the new engine:

Capability	Detail
Interactive charts	Charts support hover tooltips, click-to-filter, and direct export to image or CSV
Advanced filtering	Filters can be applied within a generated report without re-running the query
Report preview	A preview renders before the full report is committed, reducing wasted generation time
Download queue	Reports are queued and generated in the background; the UI remains usable during generation
Tag-based discovery	Reports are tagged (e.g., "Nutanix AHV", "Microsoft 365") for filtered catalog browsing
PostgreSQL backend	Report queries execute against the new PostgreSQL reporting database for improved performance

Table 11.2: New reporting engine enhancements

Export format changes

Administrators upgrading from earlier versions of Veeam ONE should be aware of export format changes introduced with the new engine:

- XML, Word (`.docx`), and Excel (`.xlsx`) export formats have been removed.
- PDF export replaces Word for formatted report output.
- CSV export replaces XML and Excel for data extraction.
- Patch 1 (v13.0.1) enhances CSV exports to include a separate Summary CSV and one CSV per chart, all bundled in a single ZIP archive.

Note

Migration note:

Organizations with automated report delivery processes that rely on XML or Excel exports will need to update their scripts or downstream processes before upgrading to v13.

Redesigned report catalog

Alongside the new engine, the **Report Catalog** in the web console has been redesigned from a flat list view to a tile/thumbnail layout. Visual preview thumbnails show an approximation of each report's layout before it is opened. Reports are filterable by purpose, tag, or a "*Popular Reports*" category. A global search covers report names and descriptions, making it significantly faster to locate the right report in large Veeam ONE deployments with many custom report templates.

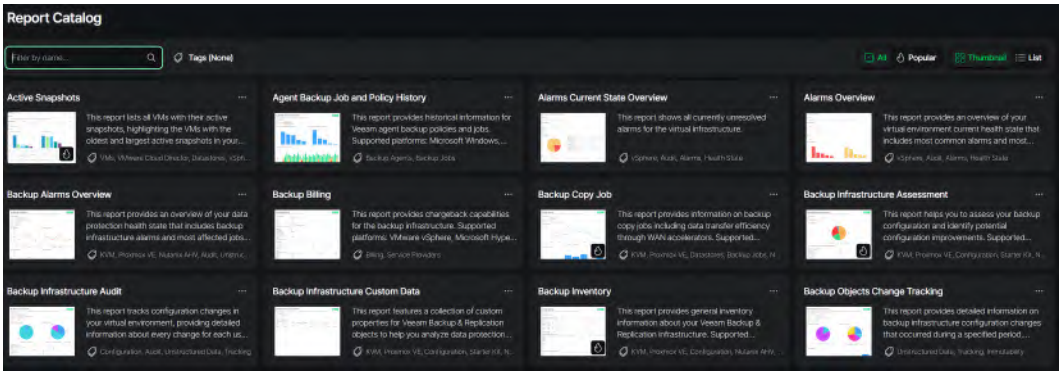


Figure 11.3: New report catalog layout

Notable report enhancements (v13.0.1) include the following:

- Now includes standalone Veeam plug-in data for enterprise applications. Workload Protection History Report.
- Now includes standalone Veeam plug-in workloads. Backup Inventory Report.
- Restores a date-as-columns heat-map view showing a week of protection outcomes per VM, with a toggle between the heat-map and row-based views. VM Daily Protection Status.

Technet24

- Optional table of contents restored in PDF exports via a user-controlled checkbox. Table of Contents for PDFs.
- Eliminates unnecessary whitespace and forced page breaks, reducing page count while preserving readability. Enhanced PDF layout.

Infrastructure monitoring updates

Veeam ONE v13 and v13.0.1 extend platform support and expand monitoring coverage across virtualization and cloud infrastructure.

Platform	Change in v13 / v13.0.1	Notes
VMware vSphere 9.0	Added support (v13.0.1)	Full monitoring, alarms, and reporting for vSphere 9.0 environments
VMware Cloud Director 10.6	Added support (v13.0.1)	Extends VCD monitoring to the latest release
VMware vSphere 6.x	Removed support	Administrators must upgrade vSphere before upgrading Veeam ONE to v13
Microsoft Hyper-V/ Windows Server 2012	Removed support	Windows Server 2012 Hyper-V is no longer a supported monitored platform
Nutanix AHV	New monitoring, alarms, reports	AHV VMs are now visible in dashboards and reports; filter Report Catalog by "Nutanix AHV" tag
Veeam Software Appliance (VSA)	New support	Linux-based VBR appliance is now a monitored infrastructure component
Microsoft Hyper-V (current versions)	Continued support	No changes to existing Hyper-V monitoring capabilities

Table 11.3: Monitoring updates

Enterprise application plug-in monitoring (v13.0.1)

One of the most-requested capabilities in v13.0.1 is monitoring support for standalone Veeam plug-ins for enterprise applications. In versions before v13.0.1, these plug-ins operated entirely outside Veeam ONE's visibility. Their backup jobs did not appear on dashboards, their workloads were excluded from protection reports, and no alarms could be raised against their policy state.

With v13.0.1, Veeam ONE gains the following capabilities for standalone plug-in workloads:

- Workload Protection History Report now includes plug-in backup data.
- Backup Inventory Report now includes plug-in workloads.
- Application Backup Policy State alarm tracks per-plugin protection compliance.
- Desktop Client displays standalone plug-in workloads in the infrastructure tree.

This improvement is particularly significant for organizations running mission-critical workloads on SAP HANA or Oracle databases, where backup compliance has previously required separate monitoring tooling outside Veeam ONE.

Cloud workload monitoring expansion (v13.0.1)

Veeam ONE v13.0.1 expands the range of public cloud workloads that can be monitored and reported on, adding coverage for several new resource types across AWS, Azure, and Google Cloud Platform.

Cloud provider	Newly supported workload types
Amazon Web Services	AWS FSx
Amazon Web Services	AWS Redshift Clusters
Amazon Web Services	AWS DynamoDB
Amazon Web Services	AWS Redshift Serverless
Microsoft Azure	Azure Virtual Network
Microsoft Azure	Azure CosmosDB
Google Cloud Platform	Google Cloud Spanner

Technet24

Cloud provider	Newly supported workload types
REST API	Optimized cloud backup endpoints with more data and reduced execution time (v13.0.1)

Table 11.4: Cloud workloads for monitoring

These additions build on the existing cloud monitoring capabilities covering AWS EC2, RDS, EFS, and VPC; Azure Virtual Machines, SQL Database, and Azure Files; and GCP Compute Engine and Cloud SQL. Organizations with multi-cloud environments can now consolidate a wider proportion of their cloud backup reporting within a single Veeam ONE instance.

Veeam Backup & Replication monitoring improvements

With Veeam Backup & Replication v13 being released, there have been significant improvements in Veeam ONE added for monitoring the environment. Let's look at them below.

Job calendar view

A new **Jobs Calendar** view in the Veeam ONE web console provides a day, week, or month visualization of backup job schedules and their outcomes.

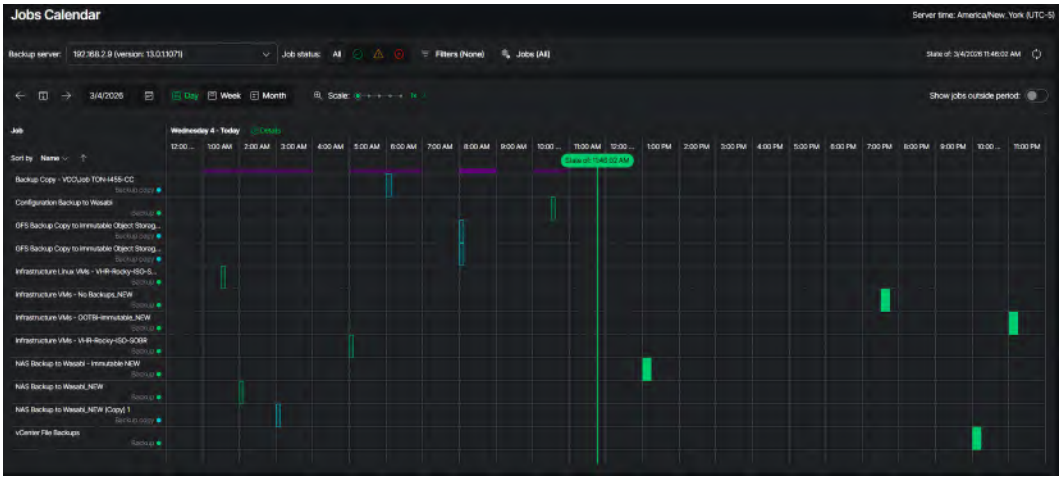


Figure 11.4: Job calendar view in Web console

Administrators can filter the calendar by job status or by backup server, making it easy to identify scheduling conflicts, missed windows, or trends in job failures over time. The calendar supplements the existing alarm-based approach to job monitoring with a temporal perspective.

Credentials manager

A new **Credentials Manager** section in the web console provides a centralized store for credential records used to connect to monitored infrastructure.

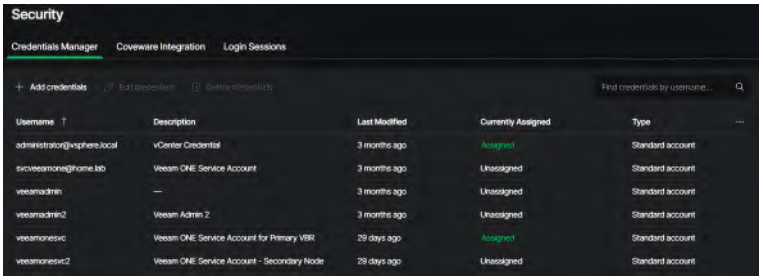


Figure 11.5: Credential manager in Web console

It supports standard Windows accounts and Linux private key (identity/public key) authentication. Veeam ONE will warn administrators if a credential record that is actively in use is being modified or deleted, preventing accidental monitoring disruptions.

Data collection configuration in web client

Data collection settings which were previously accessible only through the Veeam ONE desktop client are now configurable directly from the web console.

Technet24

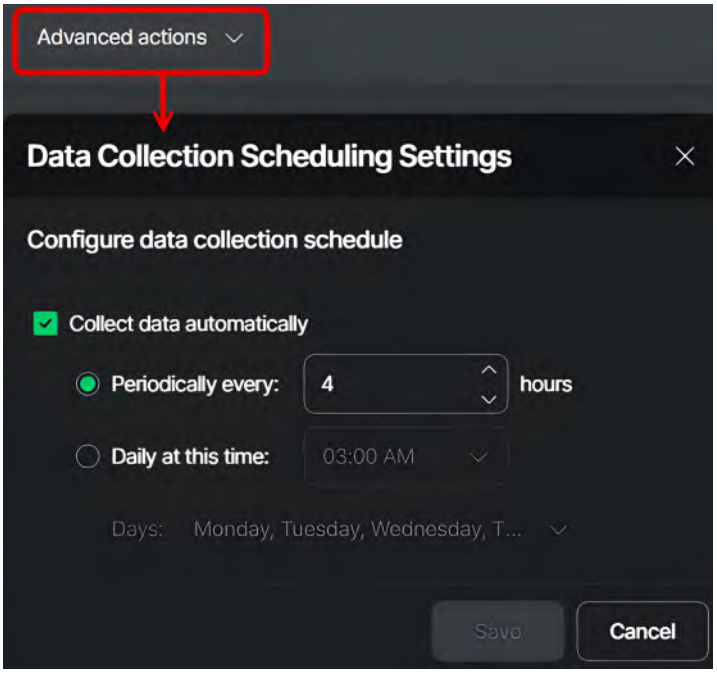


Figure 11.6: Data collection schedule settings

This aligns with Veeam's broader move toward a web-first management experience and allows administrators to manage Veeam ONE entirely from a browser without needing to install the Windows-based desktop client.

Veeam intelligent diagnostics updates

Veeam intelligent diagnostics (VID) continues to analyze Veeam Backup & Replication server logs locally, matching patterns against a signature library to surface actionable alarms with built-in remediation recommendations. In v13, signature updates are released independently of the product version and can be applied automatically. All log analysis occurs locally, and no log data is transmitted to Veeam servers.

Veeam intelligence integration

Veeam intelligence received major capability additions in v13 and v13.0.1. While not exclusively a monitoring feature, Veeam intelligence draws heavily on Veeam ONE's monitoring data and significantly changes how administrators interact with that data.

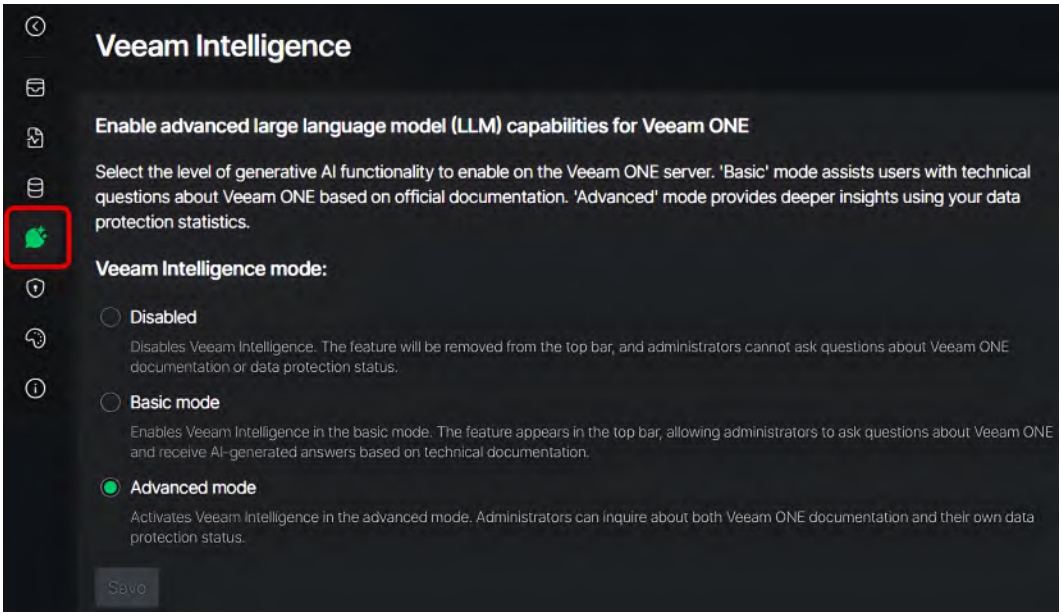


Figure 11.7: Veeam intelligence integration and settings

Malware threat analysis agent (v13.0.1)

A dedicated malware threat analysis agent within Veeam intelligence provides real-time insights into malware detections surfaced in the threat center. It can explain anomalies, estimate blast radius per workload, and provide next-best-action guidance for recovery steps. Administrators can trigger signature-based malware scans directly from the intelligence interface, shortening the time between detection and investigation.

Deep-data analysis agent (v13.0.1)

The deep-data analysis agent allows natural language queries against any Veeam ONE saved report. It processes millions of rows without requiring the report to be regenerated, loading results into an internal data frame for instant filtering, compact table generation, and on-demand chart rendering. This makes large reports such as multi-month protection history across thousands of VMs practically queryable in seconds.

Data resilience summary (v13.0.1)

The **data resilience summary**, also referred to as the **morning coffee report**, is a daily AI-generated briefing that aggregates backup job outcomes across the entire managed estate. It groups errors by workload for faster triage, adds context-aware recommendations linked to Veeam knowledge base (KB) articles, and produces an executive-level summary suitable for daily operations reviews. This requires VBR or VSA version 13.0.1 or later.

Technet24

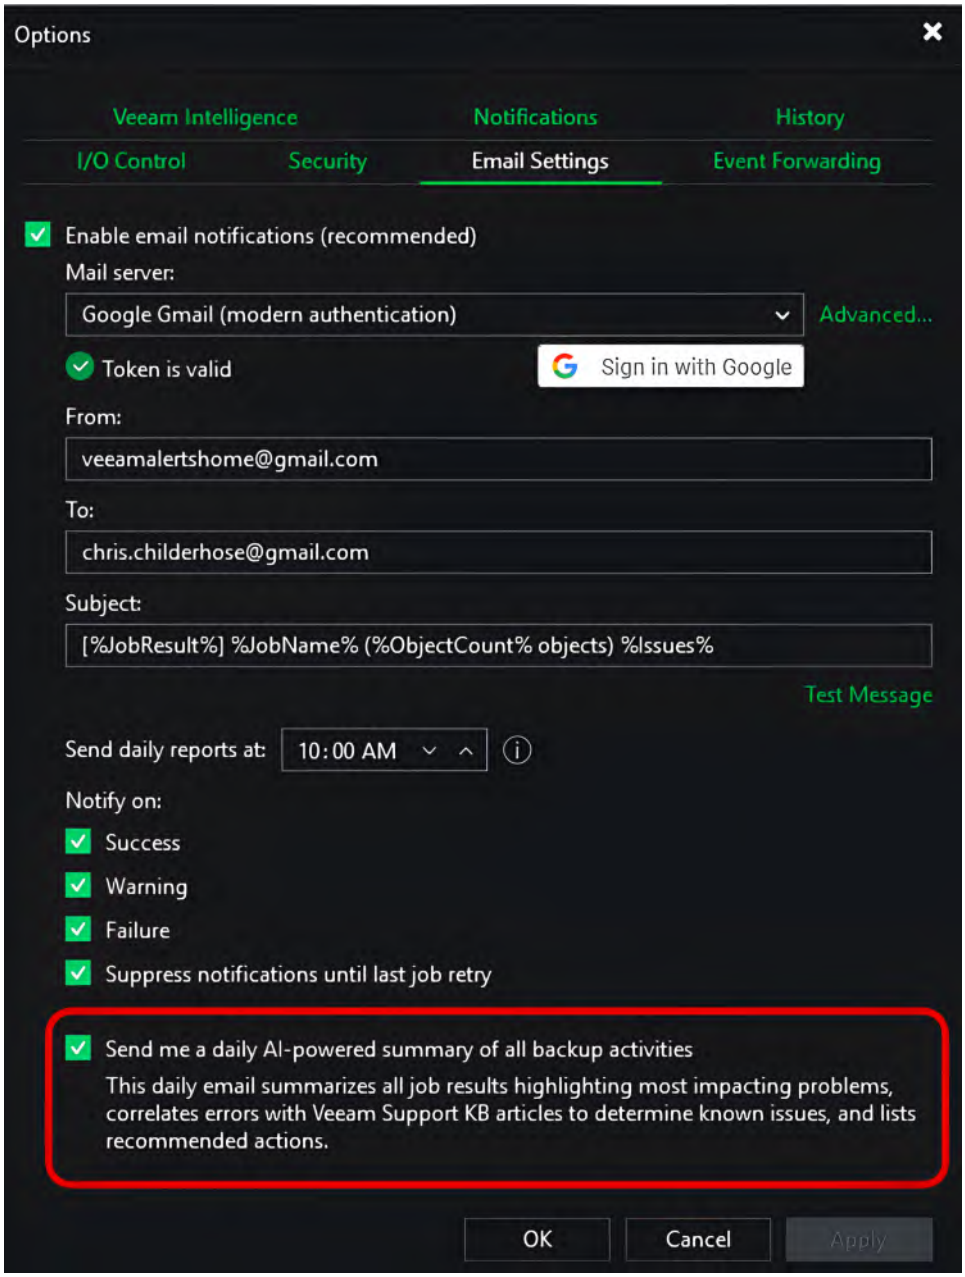


Figure 11.8: Data resilience summary option in email Settings

Other additional intelligence enhancements include the following:

- Voice input (speech-to-text) supports asking questions in natural language by microphone, with text-to-speech responses.
- Charts (donut, pie, bar, column, line) can be rendered directly within the Intelligence chat interface from monitoring data sources.
- Advanced mode exposes environment context including backup server versions, licensing, proxy configurations, repository capacity, and malware events.
- Redundant LLM instances provide failback between secure Azure OpenAI endpoints for high availability.
- Grid export options allow downloading Intelligence-generated grids as CSV or JSON.

Veeam ONE v13 represents a generational step forward in the monitoring and analytics capabilities of the Veeam data platform. The changes in this release span architecture, user experience, alarm coverage, reporting, and AI-assisted analysis. The following table summarizes the key themes and their operational impact:

Theme	Operational impact
Security-first monitoring	Threat center as the default view keeps cyber resilience front and center for daily operations.
Architecture modernization	Veeam Analytics Service and dual-database model improve scalability and reporting performance.
Reporting engine replacement	Interactive HTML5 reports replace SSRS; administrators must update export-dependent workflows.
Expanded workload coverage	Enterprise application plug-ins, new cloud services, and AHV VMs are now fully monitored.
New alarm categories	HA health, authentication anomaly, Coveware Recon, and plug-in policy alarms close significant coverage gaps.

Technet24

Theme	Operational impact
AI-powered operations	Veeam intelligence agents accelerate malware triage, daily briefings, and large dataset analysis.
Web console parity	Credentials manager and data collection settings reduce dependency on the desktop client.

Table 11.5: Veeam ONE enhancement summary

Administrators upgrading to Veeam ONE v13 should prioritize three actions:

- Reviewing and updating any automated report delivery processes that relied on XML or Excel exports.
- Verifying that all monitored VBR servers are running v12 or later before migration.
- Enabling the new threat center alarms, particularly authentication anomaly and application backup policy state, to take immediate advantage of the expanded coverage.

With these steps in place, the monitoring improvements in v13 deliver measurable operational value from the first day of deployment. Now let's take a deep dive into the architectural changes within the Veeam ONE 13 setup.

Veeam ONE v13 - Architecture changes

Veeam ONE v13 is not simply a feature update, but it represents a substantial re-engineering of the product's internal architecture. Where previous major versions evolved the monitoring and reporting capabilities within the same fundamental design, v13 changes how components communicate, how data is stored, how reports are generated, and how the product integrates with Veeam Backup & Replication servers. Several of these changes have direct operational consequences for administrators, including a reversal of network communication direction, a new database technology, and mandatory firewall rule updates.

Understanding the architectural changes in v13 is essential for three reasons. First, planning an upgrade from v12 requires knowing what will be preserved automatically, what must be reconfigured, and which firewall rules need to change before the upgrade can succeed. Second, designing a new v13 deployment requires selecting and sizing two databases rather than one, and understanding how the new analytics service changes connectivity requirements. Third, troubleshooting connectivity or data collection issues in v13 requires understanding the new

communication model as the old mental model of Veeam ONE polling its agents no longer applies.

In this section, we will look at each architectural change in detail, covering the Veeam analytics service, the dual-database model, system requirements, network ports, the new reporting engine, security changes, REST API updates, integration architecture, and the upgrade path from v12.

> **Note**
>
> **Scope:**
>
> This section covers Veeam ONE 13.0 (September 2025, build 13.0.0.5630) and Veeam ONE 13.0.1 (November 2025, build 13.0.1.5924). Changes specific to 13.0.1 are identified throughout.

Veeam analytics service

The most architecturally significant change in Veeam ONE v13 is the replacement of the legacy Veeam ONE agent with the Veeam analytics service. The component has been substantially redesigned, and the change in communication direction alone requires firewall rule updates in every environment before an upgrade can function correctly.

Veeam analytics service architecture

In v13, the Veeam analytics service operates in two distinct modes that together replace the old agent model:

- **Server/analysis mode**: Runs on the Veeam ONE Server machine itself. It's responsible for analyzing collected logs and data, managing Intelligent Diagnostics signature updates, and processing remediation logic received from connected VBR servers.
- **Client/agent mode**: Runs on each connected VBR server. It's responsible for local data collection, log forwarding, and executing remediation commands issued by the Veeam ONE server.

The critical architectural difference is the direction of communication. In v13, the client/agent mode instance running on the VBR server initiates the connection inbound to the Veeam ONE server on TCP port 2805. This is the reverse of v12 behavior. The Veeam ONE server no longer reaches out to VBR servers to collect data. VBR servers push data to Veeam ONE.

Technet24

	v12 behavior	v13 behavior
Communication initiator	Veeam ONE Server	VBR Server (Analytics Service client)
Direction	Veeam ONE → VBR server on TCP 2805	VBR server → Veeam ONE on TCP 2805
Port to open on Veeam ONE firewall	None (Veeam ONE is the initiator)	TCP 2805 inbound from all connected VBR servers
Port to open on VBR firewall	TCP 2805 inbound from Veeam ONE	None (VBR is the initiator)

Table 11.6: Analytics Service Architecture

Note

Important:

The direction reversal on TCP 2805 is the single most common cause of post-upgrade connectivity failures. Before upgrading, update firewall rules to allow TCP 2805 inbound on the Veeam ONE server from all VBR servers, and remove (or do not rely on) the old rule permitting Veeam ONE to reach VBR servers on 2805.

Deployment and push installation

The Veeam analytics service is deployed automatically to the Veeam ONE server during product installation. For VBR servers added to monitoring, Veeam ONE attempts to push the client/agent mode component automatically. The push mechanism differs by VBR version:

VBR version	Push protocol	Ports required for push
VBR 13.0.1 or later (Windows)	HTTPS via gRPC	TCP 443
VBR 12.x (Windows)	WMI + SMB (legacy)	TCP 135, 445
VBR (Linux)	SSH / HTTPS	TCP 443
Enterprise Manager 13.0.1+	HTTPS / gRPC	TCP 443, 50001

VBR version	Push protocol	Ports required for push
Enterprise Manager v12.x	WMI + SMB (legacy)	TCP 135, 445

Table 11.7: Deployment and protocols/ports required

If the automatic push fails, for example because the required ports are blocked, Veeam ONE falls back to legacy remote WMI data collection (**TCP 135** plus the dynamic port range **49152–65535**). This fallback maintains basic monitoring but is not the intended operating model and should be treated as a signal that the push connectivity issue needs to be resolved. Manual installation packages in `.msi` (Windows) and `.bndl` (Linux) formats are available from the web console under **Configuration | Data Collection | Veeam Analytics Service** for environments where push cannot be used.

These are a few additional operational notes to consider:

- **DNS requirements**: Bidirectional DNS resolution is mandatory. The Veeam ONE server must be able to resolve VBR server names via reverse DNS (PTR records), and each VBR server must resolve the Veeam ONE Server's FQDN.
- **Single-owner model**: Each analytics service instance is owned by one Veeam ONE server. Adding a VBR server to a second Veeam ONE instance overwrites ownership and disconnects it from the first.
- **High Availability clusters**: The analytics service must be installed on every node in a VBR HA cluster. Data collection occurs from the active node only.
- **Compatibility**: The Veeam analytics service requires Veeam Backup & Replication v12 or later. Environments running VBR v11 or earlier cannot connect to Veeam ONE v13 using this service.

These are some of the more important architectural changes within the Veeam ONE architecture. Now we will look at the new dual-database architecture that uses Microsoft SQL and PostgreSQL.

Dual-database architecture

Veeam ONE v12 and all earlier versions relied on a single Microsoft SQL server database for all product functions including monitoring state, alarm history, infrastructure topology, and report generation all competed for the same database resources. In high-scale environments or during complex report generation, this created performance contention that could cause monitoring delays or slow report output.

Veeam ONE v13 resolves this by introducing a dual-database model, separating monitoring and reporting workloads onto purpose-built database platforms.

Technet24

Monitoring database (SQL server)

The Veeam ONE monitoring database remains on Microsoft SQL server and continues to serve as the operational data store. It holds live monitoring data, alarm states, infrastructure topology, session history, and all time-series performance data collected from the monitored environment. This database must remain highly available because the monitoring engine reads from and writes to it continuously.

The following are the requirements for Veeam ONE monitoring database on SQL:

- **Supported versions**: SQL Server 2016 through 2025; Azure SQL Database; Amazon RDS for SQL server.
- **Minimum compatibility level**: 130 (SQL Server 2016).
- **Bundled option**: SQL Server 2017 Express is included in the Veeam ONE installer for smaller deployments, subject to the Express edition 10 GB database size cap.
- **Communication port**: TCP 1433.
- **Sizing impact**: Because heavy report generation has moved to PostgreSQL, the SQL Server instance can be sized more conservatively than in v12. Veeam provides a database calculator tool for accurate sizing guidance.

Reporting database (PostgreSQL)

The Veeam ONE reporting database is a new addition in v13. Built on PostgreSQL, it is used exclusively by the new reporting engine to store cached report data, temporary report tables, and intermediate datasets. Offloading report generation to PostgreSQL removes the largest single source of SQL server load that existed in v12.

The following are the particulars for Veeam ONE reporting dayabase:

- **Bundled version**: PostgreSQL 17.6 (as shipped with v13.0.1).
- **Supported versions**: PostgreSQL 14.x, 15.x, 16.x, 17.x, and 18.x. UTF-8 encoding is required.
- **Deployment**: Can use the co-installed instance deployed by the Veeam ONE installer, or a remote PostgreSQL server.
- **Communication port**: TCP 5432.
- **Temporary data**: Report tables are deleted automatically when a Reporter session ends or after a session timeout. The database does not grow unboundedly with report history.

Note

Known issue — Turkish locale:

> PostgreSQL 17 installation fails on Windows systems configured with the tr-TR (Turkish) locale due to non-ASCII characters in the locale name. Veeam acknowledges this in the v13.0.1 release notes. Workaround: use a remote PostgreSQL instance on a system with a compatible locale.

Implications for administrators

The dual-database model introduces new operational considerations that did not exist in v12:

- **Backup strategy**: Both databases must be included in backup plans. The SQL server monitoring database should be backed up with transaction log backups for point-in-time recovery. The PostgreSQL reporting database can be backed up with standard PostgreSQL tooling (`pg_dump` or `pg_basebackup`) though its contents are regenerable from source data.
- **Migration after upgrade**: If the monitoring database is subsequently migrated to a different SQL server instance, a manual data recollection must be triggered for the reporting database to rebuild its cached datasets.
- **Connection limit**: A known issue in both v13 and v13.0.1 causes the installer to fail if the selected PostgreSQL instance has already reached its maximum connection limit. Check `pg_stat_activity` before running the installer against an existing PostgreSQL instance.

System requirements

Veeam ONE v13 updates system requirements in several areas. The following applies to the Veeam ONE server, web services, and client components.

Operating system

All installations require a 64-bit operating system. Supported platforms are listed in *Table 11.8*:

Operating system	Notes
Windows Server 2025	Added in v13.0.1
Windows Server 2022	Full support
Windows Server 2019	Full support
Windows Server 2016	Full support

Technet24

Operating system	Notes
Windows 11 (Pro/Enterprise, 22H2 minimum)	Development and lab use
Windows 10 (1909 and LTS builds)	Development and lab use

Table 11.8: Operating systems supported

Hardware (All-in-one deployment baseline)

Hardware support for all-in-one deployment. System Resources are listed in *Table 11.9*:

Resource	Minimum	Scaling rule
CPU	4 cores (x86-64)	+0.03 vCPU per connected VBR server
RAM	8 GB (16 GB recommended)	+60 MB per connected VBR server
Disk	50 GB for product files	Additional space required for SQL Server and databases — use Veeam Database Calculator

Table 11.9: Hardware requirements

Software prerequisites

All installations require .NET and some other windows components. Supported .NET and components are listed in *Table 11.10*:

Component	Required version
Microsoft ASP.NET Core Shared Framework	8.0.17
.NET Runtime	8.0.17
.NET Framework	4.7.2
Windows Management Instrumentation (WMI)	Enabled (required service)
File and Print Sharing	Enabled (required service)

Component	Required version
Web browser (for web client)	JavaScript must be enabled

Table 11.10: Software prerequisites

New monitored platform support added in v13 / v13.0.1

Platform support has been updated. Supported platforms are listed in *Table 11.11*:

Platform	Version	Added in
VMware vSphere / ESXi	9.0	v13.0.1
Microsoft Hyper-V/Windows Server	2025	v13.0.1
Microsoft SCVMM	2025	v13.0.1
Microsoft SQL Server (monitored)	2025	v13.0.1
Veeam Backup & Replication	v12 – v13	v13.0
VMware vSphere 6.x	All	Removed in v13
Windows Server 2012 Hyper-V	All	Removed in v13

Table 11.11: New platform support for monitoring

Hopefully, each of the preceding charts helps outline the system requirements needed for installing Veeam ONE 13. Now we will look at network port requirements that are needed should you have a secure or firewalled network.

Network port reference

The v13 port map differs from v12 in two important ways: the communication direction on TCP 2805 is reversed (VBR now initiates to Veeam ONE), and new ports appear for the PostgreSQL Reporting Database (TCP 5432) and for gRPC-based push of the analytics service to VBR v13 servers (TCP 443). *Table 11.12* covers the primary communication paths:

Technet24

Source	Destination	Port	Purpose
VBR server (any version)	Veeam ONE Server	TCP 2805	Analytics Service: data push to Veeam ONE (direction reversed vs. v12)
VBR / Enterprise Manager	Veeam ONE Server	TCP 1239, 2741	Analytics tab in VBR/EM console
Veeam ONE	VBR v13+ (Windows)	TCP 443	Analytics Service push install via gRPC/HTTPS
Veeam ONE	VBR v12.x (Windows)	TCP 135, 445	WMI + SMB for Analytics Service push install
Veeam ONE	VBR (Linux)	TCP 443	Analytics Service push install
Veeam ONE	Enterprise Manager v13+	TCP 443, 50001	Remote Analytics Service installation
Veeam ONE	Backup proxy/repository	TCP 135, 49152–65535	WMI performance data (legacy collection path)
Veeam ONE	ESXi host	TCP 443	Data collection via HTTPS
Veeam ONE	ESXi host	TCP 5989	CIM XML hardware detail collection
Veeam ONE	vCenter Server	TCP 443	Data collection via HTTPS
Veeam ONE	VMware Cloud Director	TCP 443	REST API
Veeam ONE	Hyper-V host	TCP 135, 445, 49152–65535	WMI + remote registry
Veeam ONE	VM Guest (Windows)	TCP 135, 445, 49152–65535	Guest OS process/service monitoring

Source	Destination	Port	Purpose
Veeam ONE	VM Guest (Linux)	TCP 22	SSH process/service monitoring
Veeam ONE	Veeam Backup for M365	TCP 4443	REST API data collection
Veeam ONE	Veeam Backup for M365	TCP 5985, 5986	Remote service enablement + certificate install
Veeam ONE	SQL Server (Monitoring DB)	TCP 1433	Monitoring Database connection
Veeam ONE	PostgreSQL (Reporting DB)	TCP 5432	Reporting Database connection (new in v13)
Veeam ONE	Veeam intelligence (cloud)	TCP 443	rest-ai.veeam.com, cdn.rest-ai.veeam.com, js.veeam.com
Veeam ONE	License/Signature server	TCP 443, 80	License and Intelligent Diagnostics signature updates
Veeam ONE Web Service	Veeam ONE Server	TCP 2741, 2742	Internal Web API + Reporting Service
Web browser	Veeam ONE Web Services	TCP 1239	HTTPS access to web console

Table 11.12: Port reference for firewall rules

Note

Key changes are marked:

TCP 2805 direction reversal (VBR now initiates to Veeam ONE) and TCP 5432 for the new PostgreSQL reporting database are the two new firewall requirements that affect every v13 deployment.

Technet24

Reporting engine architecture

Veeam ONE v12 built its reporting on Microsoft SQL server reporting services (SSRS), a Windows-only, server-side rendering technology with limited interactivity and a fixed export pipeline. SSRS generated static documents and imposed its own service dependencies alongside the SQL server installation. In v13, Veeam replaces the Veeam-specific report set with a new HTML5/JavaScript-based interactive reporting engine, while retaining SSRS for any non-Veeam infrastructure reports that have not yet been migrated.

New engine capabilities include the following:

- **Interactive rendering**: Charts support hover tooltips, click-to-filter interactions, and export to image or CSV directly from the chart surface without re-running the underlying query.
- **In-report filtering**: Filters can be applied and changed within a generated report without triggering full regeneration, allowing rapid slicing of large datasets.
- **Preview mode**: A report preview renders with a sample of data before the administrator commits to generating the full report, avoiding wasted generation time for misconfigured reports.
- **Background download queue**: Reports are submitted to a background queue and generated asynchronously. The web console remains fully usable during generation, and completed reports appear in the queue for download.
- **Tag-based discovery**: Reports carry tags (for example "Nutanix AHV" or "Microsoft 365") that enable filtered browsing of the report catalog.
- **Visual catalog**: Thumbnail previews, purpose filters, and a "Popular Reports" category replace the flat list of the v12 catalog.

Export format changes

The new engine does not support all export formats that SSRS provides. Administrators with automated report delivery workflows should review these changes before upgrading:

Format	v12 status	v13 status
PDF	Available	Available — primary formatted output
CSV	Available	Available — primary data extract format; enhanced in v13.0.1 to include summary CSV and per-chart CSVs in a ZIP archive
XML	Available	Removed

Format	v12 status	v13 status
Word (DOCX)	Available	Removed — use PDF
Excel (XLSX)	Available	Removed — use CSV

Table 11.13: Report export formats

Data flow

When a report is requested, the new reporting engine queries the PostgreSQL reporting database rather than SQL server. The PostgreSQL database may already hold cached data from a previous run of the same report; if so, cached data is used to accelerate delivery. If no cache exists, the engine queries source data, renders the report, and caches the intermediate dataset in PostgreSQL for future use. Temporary tables are removed automatically at session end, keeping the reporting database lean.

Note

Compatibility note:

SSRS is not fully removed in v13. Reports for infrastructure types not yet migrated to the new engine continue to use SSRS. Administrators do not need to take any action—the product selects the correct engine automatically based on report type.

Web client and desktop client

Veeam ONE has shipped with both a Windows desktop client and a web-based console for several versions. In v13, the web client receives substantial new capabilities including features previously exclusive to the desktop client, which reflect Veeam's trajectory toward a fully web-based management experience. The desktop client has not been deprecated and remains required for a small number of tasks, but the web client is now the primary interface for day-to-day operations.

Capability	Web client	Desktop client
Threat center dashboard	Yes (default view)	No
Redesigned report catalog with thumbnails	Yes	No
Dark mode and theme selection	Yes	No

Technet24

Capability	Web client	Desktop client
Credentials manager	Yes (new in v13)	No
Data collection configuration	Yes (new in v13)	Yes
Alarm configuration	Yes	Yes
Dashboard editing	Yes	Yes
Report generation and scheduling	Yes	Yes
Veeam intelligence (AI assistant)	Yes	No
Job calendar view	Yes (new in v13)	No
Enterprise application plug-in visibility	Yes (v13.0.1)	Yes (v13.0.1)
ReFS repository credential configuration	No	Yes (desktop only)
Build version shown after v13.0.1 patch	13.0.1.5924	13.0.0.5630 (unchanged)

Table 11.14: Webclient and desktop client options

Note

Practical note:

After applying the v13.0.1 patch, the desktop client version number does not change — it remains at 13.0.0.5630. Only the web client builds number updates to 13.0.1.5924. This is by design and does not indicate an incomplete patch application.

Security architecture changes

With architecture changes and enhancements, we cannot forget security. We'll now discuss some of the many security changes and enhancements within Veeam ONE 13.

Protocol modernization: gRPC replaces WMI/RPC

For components running VBR v13, Microsoft RPC and WMI are no longer used for inter-component communication. They are replaced by gRPC, which is a modern, cross-platform remote procedure call framework running over HTTP/2. This change reduces the number of ports required for push installation of the analytics service to a single port (TCP 443) and

improves communication efficiency and security through HTTP/2's built-in TLS. For VBR v12.x servers, the legacy WMI/RPC path remains in use.

NTLM deprecation

NTLM authentication between Veeam ONE components is deprecated in v13 in favor of Kerberos. Environments that have disabled Kerberos or rely on NTLM-only domain configurations should verify compatibility before upgrading. The Veeam Deployment Kit, installable on Windows infrastructure components, provides certificate-based authentication as an alternative where Kerberos is not available.

SAML-based single sign-on

Veeam ONE v13 introduces support for SAML-based single sign-on (SSO) with external identity providers. Microsoft Entra ID (formerly Azure Active Directory), Okta, and other SAML-compatible providers are supported. When SSO is configured, MFA can be enforced through the identity provider, removing the need for Veeam ONE to manage MFA policies directly.

RBAC and role changes

The role-based access control model in v13 introduces new built-in roles alongside enhanced custom role capabilities:

Role	Purpose	Notes
Incident API operator	Service account role for REST API access to the Incident API endpoint only	No VBR console access; MFA not supported; disabled by default
Security administrator	Dedicated administrative role with security configuration rights	New in v13
Custom roles	Scoped permissions configurable per-role	Enhanced via role assignment wizard in web client

Table 11.15: RBAC and role changes

Technet24

REST API authentication

The Veeam ONE REST API uses OAuth 2.0 as its authentication framework. Clients obtain an access token and a refresh token. The access token is passed in the Authorization header of all subsequent API requests; when it expires, the refresh token is used to obtain a new access token without re-entering credentials. This model is compatible with modern automation and integration toolchains.

We have now covered the many security enhancements within Veeam ONE 13, so now let's look at some of the integration architectural enhancements in this release.

Integration architecture

Veeam ONE 13 introduces some integration architecture changes including Veeam intelligence, Coveware Recon Scanner, and Intelligence signature updates. We will discuss each of these topics below.

Veeam intelligence

Veeam intelligence is the AI-powered assistant integrated into Veeam ONE. Understanding its connection model is important for environments with strict egress firewall policies, because Veeam intelligence relies on cloud-hosted AI services rather than on-premises processing.

The communication path is as follows: **web browser | Veeam ONE Server | Veeam AI cloud services**. The Veeam ONE server acts as a proxy and AI queries do not leave from the browser directly to the cloud. Three outbound connections from the Veeam ONE server to TCP 443 are required:

- `rest-ai.veeam.com`: Primary AI service endpoint (HTTPS and WebSocket/WSS).
- `cdn.rest-ai.veeam.com`: CDN for front-end content delivery.
- `js.veeam.com`: Front-end JavaScript resources.

When Veeam intelligence operates in advanced mode, it issues GET requests against the Veeam ONE REST API (internal, on the same server) to retrieve environment-specific data before formulating responses. This includes alarms, protected workloads, repository capacity, job inventory, proxy configurations, and malware events. No backup data but only structured metadata about the environment is sent to the cloud.

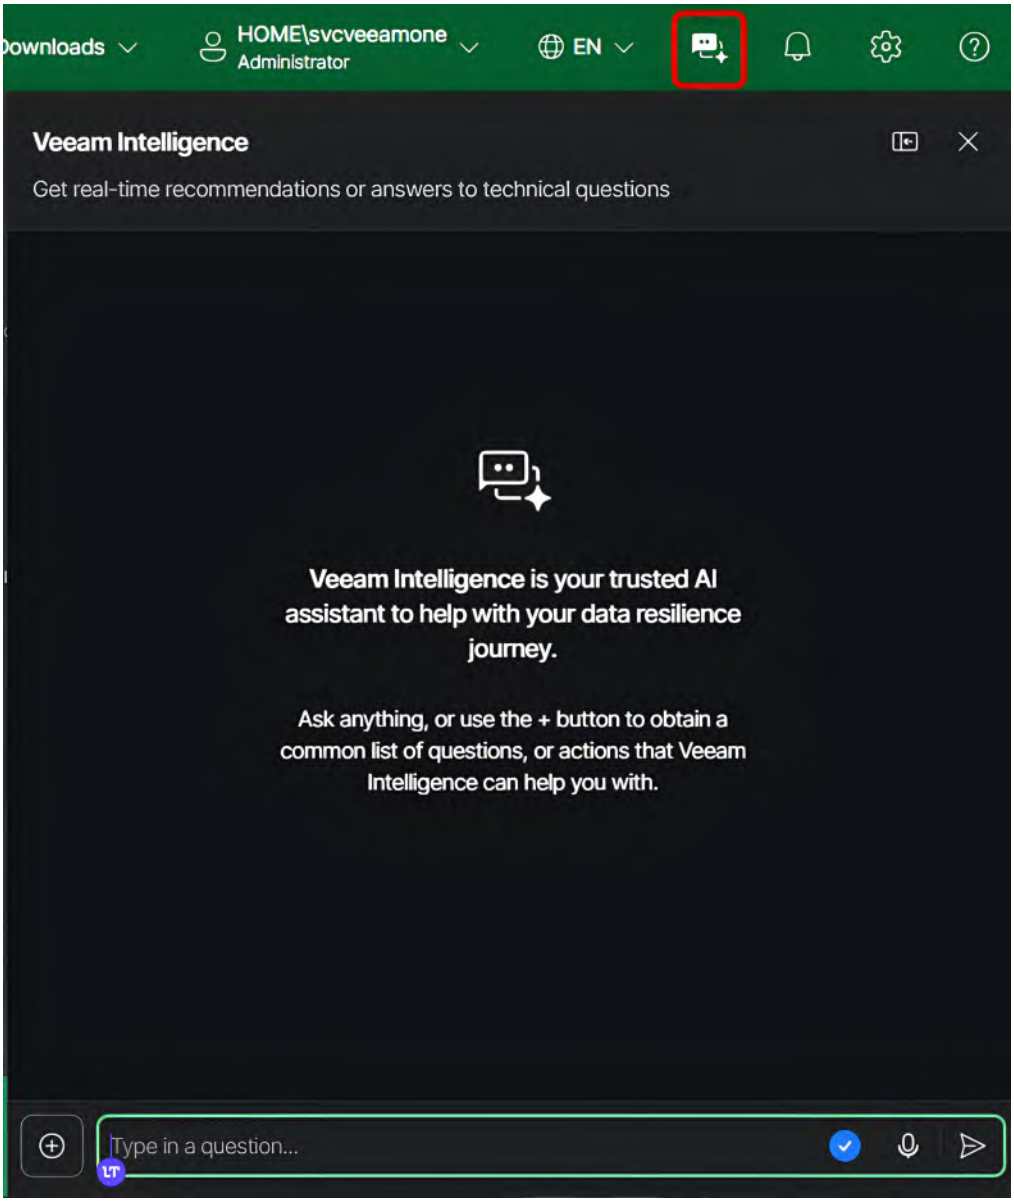


Figure 11.9: Veeam intelligence dialog within the Web console

Coveware recon scanner integration

Coveware recon scanner 3.0 is included with Veeam Data Platform premium and VCSP offerings. The integration, available in v13.0.1, surfaces recon scanner findings directly in the

Technet24

Veeam ONE threat center as actionable alarms. Each incident flagged by recon generates a corresponding Veeam ONE alarm that can be routed through standard notification channels. Additionally, recon scanner findings can be correlated with Microsoft Sentinel, mapped against the MITRE ATT&CK framework for threat intelligence enrichment.

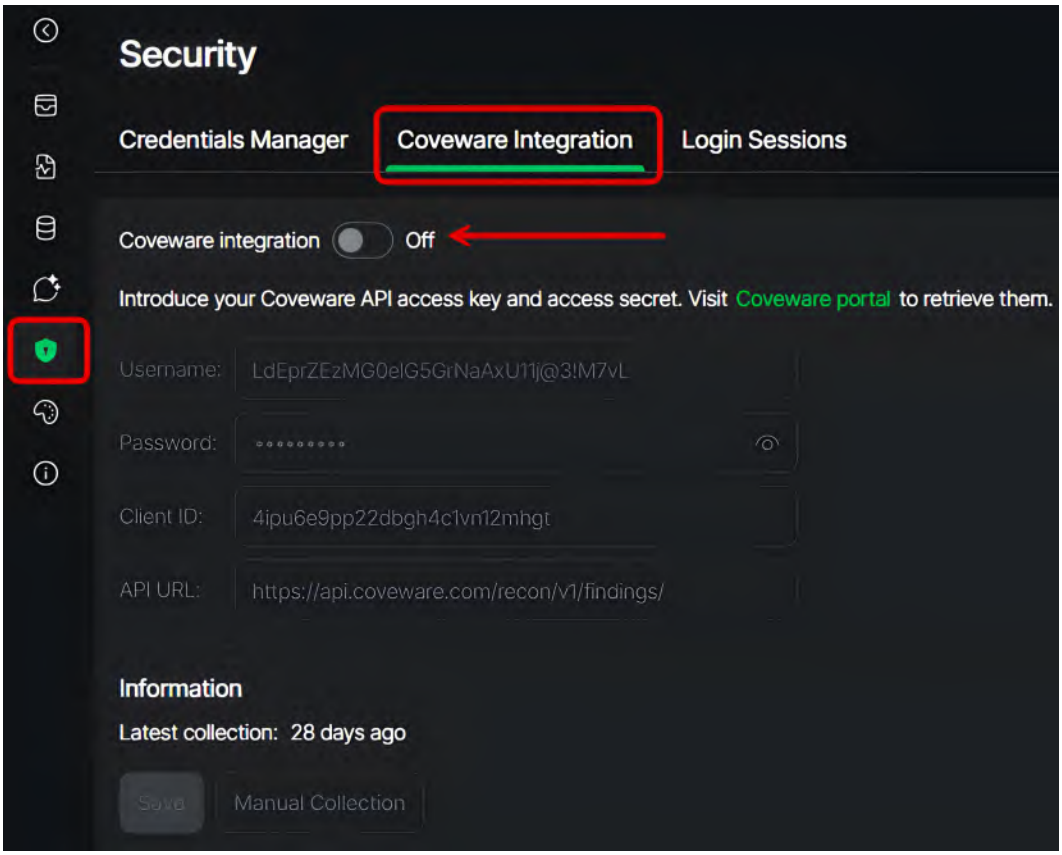


Figure 11.10: Coveware Integration in the Web console

Intelligent diagnostics signature updates

Veeam intelligent diagnostics (VID) performs on-premises log analysis against a signature library that is updated independently of product releases. Veeam ONE connects Veeam's update servers on TCP 443 or TCP 80 to download new signature files. No log data is transmitted to Veeam servers, and the analysis runs entirely locally. Signature updates can be applied automatically or held for manual approval, depending on administrator preference.

Upgrade path from v12 to v13

Upgrading from Veeam ONE v12 to v13 is supported through the in-place upgrade wizard. A full database backup before starting is mandatory as the upgrade wizard does not create one automatically. The minimum supported source version is Veeam ONE 12.0.0.2498; earlier v12 builds must first be patched to this minimum before the v13 upgrade can proceed.

Table 11.16 consolidates the pre-upgrade checklist:

Action	Reason
Back up the SQL server monitoring database	Enables rollback if the upgrade fails
Update firewall rules: open TCP 2805 inbound on Veeam ONE server from VBR servers	New communication direction — without this, Analytics Service will not connect post-upgrade
Verify all VBR servers are running v12 or later	Analytics Service is incompatible with VBR v11 and earlier
Verify bidirectional DNS resolution between Veeam ONE and VBR servers	Required for Analytics Service client/server communication
Confirm PostgreSQL instance availability (or plan for co-installed PostgreSQL)	The Reporting Database is a new mandatory component in v13
Document any custom analytics service port overrides	Known bug: custom port overrides are reset to 2805 during upgrade from v12.3
Identify reports using XML/Word/Excel export	These formats are removed in v13; downstream processes need updating
Note any custom report filters	Custom filters are not retained after upgrade and must be recreated

Table 11.16: Pre-Upgrade checklist – v12 to v13

Upgrade wizard overview

The v13 upgrade wizard guides administrators through *14* steps. *Steps 8 and 10* are new compared to the v12 upgrade process and deserve particular attention:

Technet24

Step	Action	Notes
1–4	Select components, accept licenses, review upgrade scope	Standard upgrade preparation
5	Provide updated license file	Only required if the license has changed
6	Enter service account credentials	
7	System configuration check	Installer validates prerequisites; review and resolve any failures before proceeding
8	***Configure the reporting database (PostgreSQL)***	***New step — no equivalent in v12. Specify the PostgreSQL instance and database name***
9	Specify communication ports	Defaults to TCP 2805; custom overrides reset to default here (known bug)
10	***Upgrade analytics service on connected VBR servers***	***New step — wizard pushes the Analytics Service to all reachable VBR servers; manual install required for unreachable servers***
11	Execute the upgrade	Database schema migration and service restart; duration scales with database size
12	Configure or replace default TLS certificate	Opportunity to deploy organisation-issued certificate
13	Set guest OS credentials (Hyper-V only)	
14	Apply available post-upgrade patches	

Table 11.17: Upgrade steps

The following migrates automatically:

- All existing monitoring data and historical data in the SQL Server Monitoring Database is preserved in place.

- Existing alarm configurations, thresholds, and notification settings migrate.
- Dashboard configurations and layout migration.
- Veeam-related reports are migrated to the new reporting engine automatically.
- Service account settings and connected server configurations migrate.

This requires manual reconfiguration:

- **PostgreSQL reporting database**: Must be configured during *Step 8* of the upgrade wizard. There is no equivalent in v12.
- **Firewall rules**: TCP 2805 direction must be reversed before the analytics service will function.
- **Custom report filters**: Not retained after upgrade; must be recreated manually.
- **Shared report links**: Become invalid post-upgrade.
- **Analytics service custom port overrides**: Reset to `2805` during upgrade; reconfigure after completing the wizard.
- **ReFS repository credentials**: Must be manually reconfigured via the desktop client.
- **Report export automation**: Scripts or scheduled tasks that relied on XML, Word, or Excel export must be updated to use PDF or CSV.

Note

Upgrade order:

Upgrade Veeam ONE before upgrading Veeam Backup & Replication. Veeam ONE v13 supports monitoring VBR v12 servers, so the two products do not need to be upgraded simultaneously. Upgrading VBR first would leave a v12 Veeam ONE attempting to connect to a v13 VBR server, which is an unsupported configuration.

The architectural changes in Veeam ONE v13 are broader than any previous major release. They affect not only the product's internal design but also the network configuration, database infrastructure, and operational processes of every environment that deploys it. *Table 11.18* below summarizes the key changes and their operational impact.

Area	What changed?	Operational impact
Veeam analytics service	Replaces Veeam ONE Agent; communication direction on TCP 2805 reversed	Firewall rules must be updated before upgrade; push install method changes by VBR version

Technet24

Area	What changed?	Operational impact
Dual-database model	SQL Server (monitoring) + PostgreSQL (reporting)	New database to provision, configure, and backup; SQL Server sizing can be reduced
Reporting engine	SSRS replaced with HTML5/ interactive engine	XML, Word, and Excel export formats removed; automated export workflows must be updated
gRPC adoption	WMI/RPC replaced for VBR v13 communication	Reduces ports required; improves security; WMI still used for VBR v12 and proxies/repositories
SAML SSO	Identity provider authentication supported	MFA can be enforced via IdP; reduces local credential management
Web client parity	Credentials Manager, Data Collection config moved to web	Desktop client still required only for ReFS credential configuration
Veeam intelligence	Cloud-connected AI with environment-aware Advanced Mode	Three outbound HTTPS connections required; no backup data leaves the environment
REST API v2.3	New endpoints; OAuth 2.0 authentication	Existing API integrations should be tested against v2.3 before production upgrade
OS/platform support	vSphere 9.0, WS 2025, SCVMM 2025 added; vSphere 6.x and WS 2012 Hyper-V removed	Environments running removed platforms must upgrade before moving to Veeam ONE v13

Table 11.18: Architectural changes—summary

For administrators planning an upgrade, the following are the priority actions:

- Update TCP 2805 firewall rules before the upgrade window.
- Provision a PostgreSQL instance for the reporting database.

- Audit any automated report export workflows for format compatibility.
- Verify that all VBR servers in the estate are running v12 or later.

Addressing these four areas in advance eliminates the most common causes of upgrade complications and ensures Veeam ONE v13 is operational immediately after the upgrade wizard completes. We will now look at some further enhancements to dashboards, reporting, and Veeam intelligence.

New dashboard, Veeam intelligence, and reporting changes

There have been many changes we have discussed so far with Veeam ONE v13, and we will now look at some further enhancements to things like dashboards and reporting.

> **Note**
>
> Veeam intelligence will be discussed in the section following the dashboard and reporting enhancements.

The redesigned web client experience

The Veeam ONE v13 web client has undergone a complete visual and functional overhaul. Upon logging in, administrators are greeted with the threat center, which serves as a centralized dashboard highlighting high-risk anomalies, recovery point objective (RPO) issues, and other critical insights into the environment. This provides an instant pulse check on the overall health and security posture of the backup infrastructure.

Dashboard catalog redesign

One of the most immediately noticeable changes is the redesigned dashboard catalog. Previous versions of Veeam ONE presented dashboards in a somewhat utilitarian list or grid of large tiles. In v13, the catalog has been transformed with a sleek, modern design that emphasizes visual clarity and efficient navigation. The updated layout significantly reduces the time spent searching for specific metrics and dashboards.

Key improvements to the dashboard catalog include the following:

- **Starred dashboards**: Administrators can now star their most-used dashboards, allowing them to be pinned and accessed instantly without scrolling through the entire catalog. This personalization feature is especially valuable in large environments where dozens of dashboards may be available.

Technet24

- **Visual preview tiles**: Each dashboard is now presented with a visual thumbnail preview showing the types of widgets and data visualizations it contains. This makes it much easier to identify the right dashboard at a glance without needing to open each one individually.
- **Category organization**: Dashboards are logically grouped by solution areas, including Veeam Backup & Replication, vSphere, Hyper-V, VMware Cloud Director, and security-focused categories. This organizational approach helps administrators navigate directly to relevant monitoring views.
- **Drag-and-drop customization**: Custom dashboards can now be assembled by dragging and dropping widgets from various widget packs. Combining vSphere metrics alongside backup job performance data on a single dashboard is straightforward.

Appearance and theming

Veeam ONE v13 introduces appearance customization with three modern UI themes available in both light and dark modes. The dark mode option is particularly welcome for administrators who spend extended periods monitoring their environments, as it reduces eye strain during long sessions. The theming options extend across the entire web client interface, including dashboards, reports, configuration panels, and alarm views.

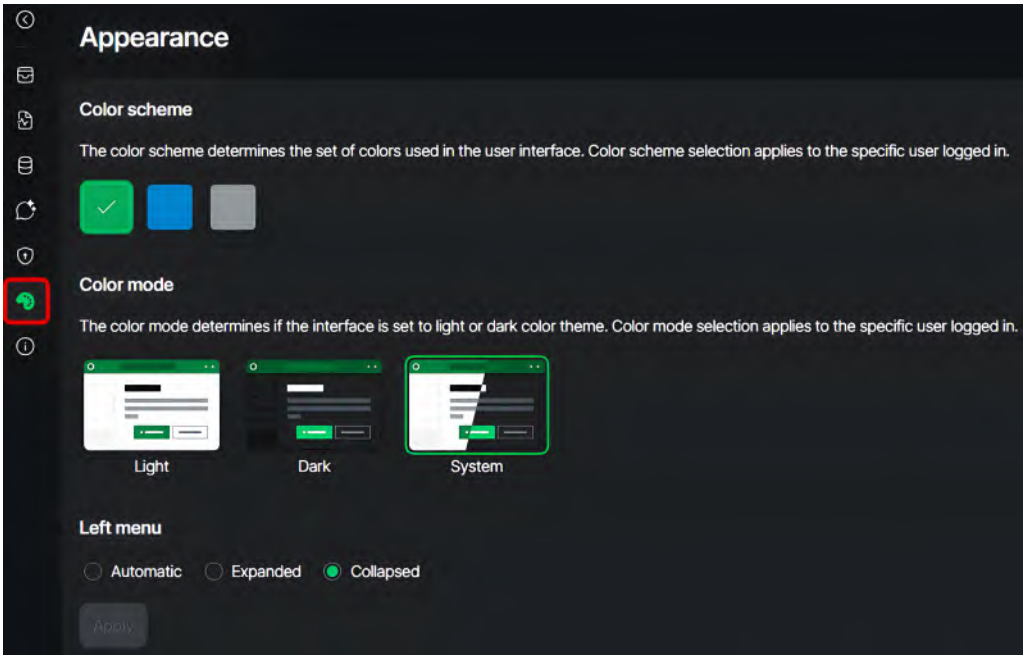


Figure 11.11: Appearance and theme for web console

Global search

A powerful global search capability has been added to the web client. Administrators can begin typing keywords from the main interface and instantly navigate relevant reports, dashboards, widgets, or alarms. The search indexes report names, descriptions, dashboard titles, and widget labels, making it possible to locate any element in the Veeam ONE environment within seconds. For example, searching for "*protected VMs*" immediately surfaces the relevant dashboard, reports, and alarm configurations related to virtual machine protection status.

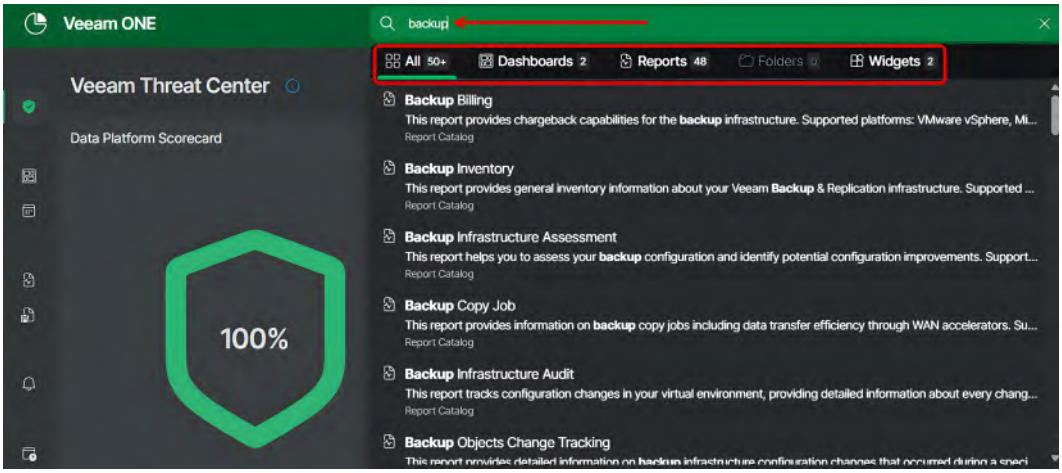


Figure 11.12: Global search from the main page of web console

Credentials manager and data collection

The web client now features a built-in credentials manager and data collection configuration section. Previously, credential management required navigating through multiple settings panels. In v13, all credentials used for connecting to Veeam Backup & Replication servers, hypervisors, and other monitored infrastructure are managed from a single, centralized location. The data collection section provides visibility into how and when data is being gathered from each monitored source.

Navigation structure

The web client navigation in v13 is organized into clearly defined sections: threat center, dashboards, jobs calendar, report catalog, saved reports, alarms overview, scheduling for reports and dashboards, deployment projects, and configuration. Each section is accessible from the main navigation pane, providing structured access to all available monitoring and reporting capabilities.

Technet24

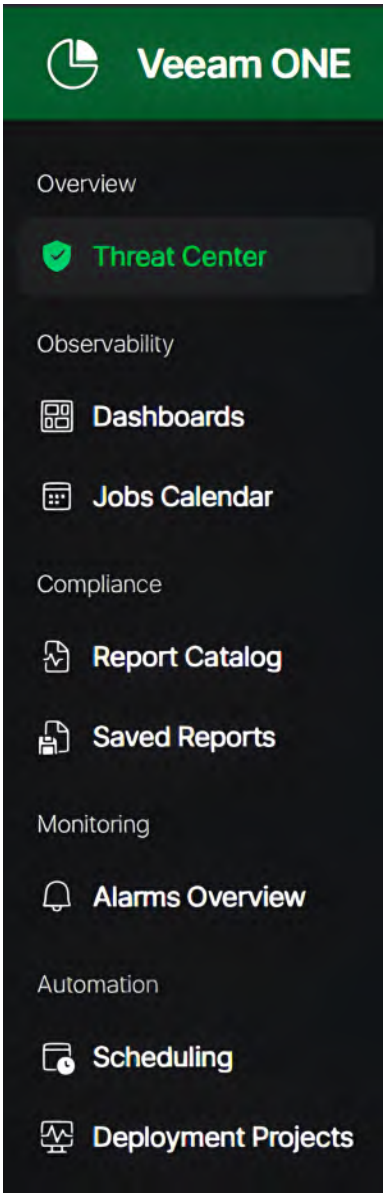


Figure 11.13: Navigation of the web console

Interactive dashboards and widgets

Dashboards in Veeam ONE v13 continue to serve as the primary at-a-glance monitoring interface, composed of widgets that display various aspects of the managed environment. Each

widget is a visual module that portrays backup infrastructure metrics in the form of interactive charts, graphs, or tables. What has changed fundamentally in v13 is the level of interactivity and the performance of these visual elements.

Predefined dashboard highlights

The Veeam Backup & Replication dashboard, one of the most used predefined dashboards, includes key widgets such as the following:

- **Backup infrastructure inventory**: Shows the types of objects in the backup infrastructure, their count, and health states.
- **Best practices findings**: Displays best practice scores by solution server, helping identify configuration drift.
- **Backup window**: Shows the total duration of daily backup and replication job sessions to identify scheduling conflicts.
- **Jobs status**: Provides an overview of successfully completed jobs versus those with warnings and errors.
- **Protected VMs overview**: Displays counters for VMs protected with backup and replication jobs.
- **Top repositories by used space**: Shows the 10 repositories closest to capacity exhaustion with trend indicators.

Heatmaps

An especially useful view available in v13 is the **Heatmaps** dashboard, accessible through the dashboards section of the web client. Heatmaps provide a time-based visualization of how backup repositories are being consumed throughout the day. Clicking on any specific period reveals detailed utilization metrics, making it possible to identify peak consumption windows and plan storage capacity accordingly. This level of temporal granularity is particularly valuable for environments with complex backup schedules spanning multiple time zones.

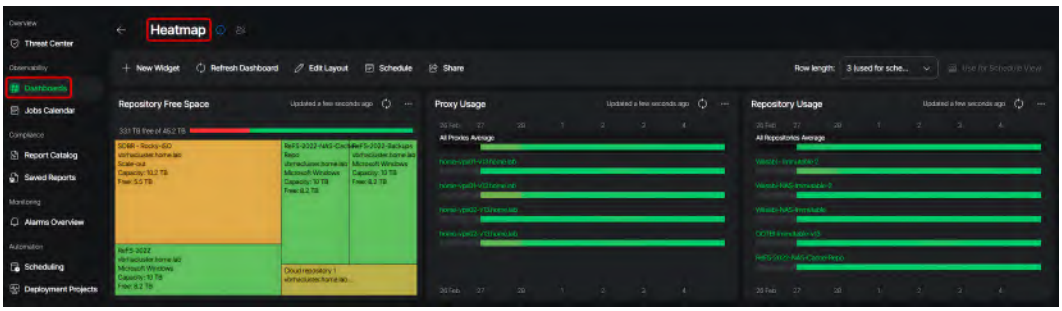


Figure 11.14: Heatmap available in the dashboard section

Technet24

Custom dashboard creation

Veeam ONE comes with widget packs organized by solution area, including packs for Veeam Backup & Replication, VMware vSphere, and Microsoft Hyper-V. When creating a custom dashboard, administrators select widgets from these packs and configure their properties including scope, time range, and display options. Widgets can be rearranged, resized, and configured independently. Each widget in a custom dashboard is linked to its underlying report, allowing administrators to click through from a summary visualization to a full detailed report with a single action.

> **Note**
>
> **Best practice:**
>
> Create role-specific custom dashboards for different teams. A security-focused dashboard combining Threat Center widgets with malware detection alerts serves security operations, while a capacity planning dashboard combining repository utilization with trend forecasting serves infrastructure teams.

Veeam intelligence: AI-powered observability

Veeam intelligence is an AI-powered solution integrated directly into the Veeam ONE web client, designed to enhance data management and decision-making through conversational interaction. First introduced in Veeam Data Platform 12.3, Veeam intelligence has been significantly expanded in v13 and especially in the 13.0.1 update with the addition of dedicated AI agents for deep analysis and actionable insights.

Architecture and integration

Veeam intelligence operates by leveraging retrieval-augmented generation (RAG) technology combined with the Veeam ONE API. This architecture allows the AI to retrieve information directly from the monitored environment, including alarms, events, threat analysis findings, backup job details, and server management data. Network communication flows between the administrator's web browser, the Veeam ONE server, and the Veeam AI service hosted at `rest-ai.veeam.com`.

Access to Veeam intelligence is through a chat button located in the upper right corner of the Veeam ONE web client window (see *Figure 11.9*). The interface supports both typed queries and, as of version 13.0.1, speech-to-text input, allowing administrators to ask questions using their voice in their own language. Responses are generated in real time, token by token, with relevant links to Veeam documentation and Knowledge Base articles integrated directly within the answers.

Operating modes

Veeam intelligence operates in two modes. The standard mode provides general guidance based on Veeam's KB and documentation. The advanced mode enables additional searches using Veeam ONE API GET queries to deliver answers specific to the administrator's own environment. In advanced mode, Veeam intelligence has access to environmental data to generate answers that reflect the actual state of the infrastructure being monitored.

> **Note**
>
> The Veeam intelligence session maintains conversation context within a session. If the chat window is closed, the context is reset and previous questions must be re-stated. Screenshots and images are not supported in the current version of the chat interface.

Malware threat analysis agent

Introduced in version 13.0.1, the malware threat analysis agent is a dedicated AI agent that provides real-time insights into security anomalies detected across the monitored environment. Through the Veeam intelligence chat interface, administrators can receive a comprehensive overview of detected anomalies, understand the blast radius of a threat, identify the most recent recommended clean restore point, and trigger actionable remediation steps.

The malware threat analysis agent can perform several key functions directly from the chat:

- **Anomaly overview**: View all detected malware anomalies organized by type and severity.
- **Blast radius assessment**: Understand which workloads and restore points are potentially affected.
- **Clean restore point identification**: Pinpoint the most recent restore point that predates the detected anomaly.
- **Ad-hoc scan triggering**: Initiate on-demand malware scans of specific workloads.
- **SureBackup verification**: Trigger SureBackup jobs to verify restore point integrity.

Deep data analysis agent

Also introduced in 13.0.1, the **deep data analysis agent** harnesses the datasets that Veeam ONE collects to allow users to be more interactive with reporting. Instead of manually navigating the report catalog, locating the correct report, configuring its parameters, and waiting for generation, administrators can simply ask Veeam intelligence to generate a report through natural language. For example, asking "*Show me the backup job success rate for the past 30 days*" will locate and generate the appropriate report without any manual navigation.

Technet24

This agent extracts reporting insights and allows administrators to quickly navigate datasets, significantly reducing the time spent working with reports. The combination of the deep data analysis agent with the new PostgreSQL-powered reporting engine means that report generation requested through Veeam intelligence is both fast and comprehensive.

Data resilience summary

Version 13.0.1 also introduces the **data resilience summary**, a consolidated snapshot of the backup environment that aggregates job status, groups errors by category, and delivers recommendations with direct links to trusted Veeam resources. This summary view, accessible through Veeam intelligence, provides a rapid assessment of overall data protection health without needing to navigate multiple dashboards or reports individually.

The new reporting engine

The reporting infrastructure in Veeam ONE v13 represents one of the most fundamental architectural changes in the product's history. Veeam has replaced the traditional SQL server reporting services (SSRS) based engine with a completely new JavaScript-based reporting engine that delivers an HTML5 experience. This migration touches every aspect of how reports are generated, rendered, interacted with, and distributed.

PostgreSQL reporting database

A cornerstone of the new reporting architecture is the introduction of a dedicated PostgreSQL reporting database. This database is purpose-built to store historical reporting data, performance metrics, and analytical datasets. It operates alongside the existing Microsoft SQL server monitoring database, which continues to handle real-time monitoring data.

The following table illustrates the Veeam ONE v13 database architecture:

Component	Database engine	Purpose	Status in v13
Monitoring database	Microsoft SQL server	Real-time monitoring data, alarms, events	Retained (existing)
Reporting database	PostgreSQL 17.x	Historical reports, cached data, warehousing	New in v13

Table 11.19: Database architecture for monitoring and reporting

It is important to understand that v13 has not replaced SQL server dependency entirely. The monitoring database continues to use Microsoft SQL server, and for the initial release, the

included SQL Server 2017 Express Edition retains its 10 GB database size limit. The PostgreSQL database specifically serves the reporting engine, enabling cached reporting, and data warehousing functionality that dramatically accelerates report generation for large environments.

> **Note**
>
> **Key architectural point:**
>
> PostgreSQL is used exclusively for the reporting database in v13. The full migration to PostgreSQL for all Veeam ONE database functions is expected in a future release. During upgrade, the installer presents a new reporting database step where administrators select either a local installation of PostgreSQL or connect to an existing PostgreSQL instance.

HTML5-based interactive reports

All Veeam-related reports have been migrated to the new HTML5-based reporting experience. Reports are no longer static PDF-style documents generated through SSRS. Instead, they are interactive web-based views that allow administrators to filter, sort, group, and rearrange columns in real time. Charts and visualizations within reports can be exported to image or CSV formats for inclusion in presentations or external analysis tools.

Key reporting improvements include the following:

- **Interactive charts and widgets**: Report visualizations respond to clicks and filters, allowing drill-down into specific data points.
- **Advanced filtering**: Reports support refined filtering for more targeted data collection and analysis.
- **Data export**: Report data can be exported to CSV and PDF formats for external distribution.
- **Background generation**: Reports download queue allows administrators to continue working while reports are generated in the background, eliminating the need to wait for large reports to complete.
- **Enhanced report builder**: The report builder has been improved with the ability to combine multiple reports into unified views with custom layouts.

Redesigned report catalog

The report catalog in the web client has been completely redesigned with a visual tile-based interface. Each report is presented with a thumbnail preview, descriptive tags, and purpose filters. A popular reports filter highlights the most used reports, helping new administrators

Technet24

discover the most valuable reporting capabilities. The catalog also supports the global search functionality, making it possible to locate specific reports by name, description, or keyword.

Veeam ONE v13 ships with over 150 prebuilt reports covering capacity planning, performance analysis, SLA tracking, trend monitoring, chargeback calculations, and security compliance. Reports are organized by category and tagged for easy discovery, with each category addressing a specific operational concern.

Table 11.20 lists the key report categories in Veeam ONE v13:

Report category	Description
Backup infrastructure	Inventory, job status, protected workloads, backup window analysis
Security and compliance	Backup compliance scores, immutability status, security hardening findings
Capacity planning	Repository utilization trends, growth forecasting, storage consumption
Performance	Job duration analysis, proxy utilization, throughput metrics
Malware detection	Anomaly reports, threat findings, malware detection summaries
Chargeback	Compute and storage consumption by customer, department, or tenant

Table 11.20: Report categories in Veeam ONE

Reporting database configuration

The PostgreSQL reporting database can be configured and managed through the web client's configuration section under database settings. Administrators can specify the PostgreSQL server hostname or IP address, database name, connection port, and authentication method. Both Windows authentication and native PostgreSQL database authentication are supported. The configuration panel also provides a test connection function, an activity audit log showing connection status and runtime information, and a cache clearing option for troubleshooting.

A retention policy can be configured to control how long historical reporting data is stored in the PostgreSQL database. This is important for managing database size in large environments where years of historical backup data could result in significant storage consumption.

Threat center integration and security reporting

The threat center, first introduced in Veeam ONE 12.1, has been significantly enhanced in v13 with deeper integration with Veeam threat hunter and indicators of compromise (IOC) capabilities. The threat center serves as the default landing page in the web client, providing an immediate security overview that includes a security scorecard, global malware detection map, RPO anomalies overview, and SLA compliance overview.

Coveware recon integration

In version 13.0.1, the threat center gains integration with Coveware recon scanner findings. recon scanner 3.0, developed by Coveware (acquired by Veeam), provides real-time operational threat visibility by flagging suspected adversary behavior across monitored endpoints, including brute force attacks, suspicious file activity, and unexpected network connections. Recon findings are now surfaced directly within the threat center with actionable alarms for each incident and filtering options to reduce alarm fatigue.

ServiceNow integration enhancements

The ServiceNow integration, which creates incidents directly within a ServiceNow instance when Veeam ONE alarms are triggered, has been enhanced in v13 with improved severity mapping. Administrators can configure the impact and urgency levels that correspond to each Veeam ONE alarm severity (error, warning, and information), ensuring that the most critical backup infrastructure issues receive appropriate priority in the ServiceNow ticketing workflow.

The bidirectional nature of the integration means that when an alarm is resolved in Veeam ONE, the corresponding ServiceNow incident is also resolved. This closed-loop approach eliminates stale tickets and provides accurate visibility into the current state of backup infrastructure issues from within the ServiceNow platform. Over 300 alarms are available in Veeam ONE v13, all of which can be configured to generate ServiceNow incidents.

Monitoring the VSA

For environments running the new VSA, monitoring through Veeam ONE v13 requires that remote data collection be enabled in the host management web UI of the appliance before adding it as a monitored server. This is a new requirement specific to the appliance deployment model, as the VSA's hardened Linux-based architecture requires explicit configuration to allow external monitoring connections.

Note

Important configuration step:

Before adding a Veeam Software Appliance to Veeam ONE v13, navigate to the Host Management Web UI on the appliance and enable the Remote Data Collection option.

Technet24

> Without this step, the Analytics Service will not be able to establish a connection for data collection.

Integration with the VBR console

One of the most notable integration improvements in v13 is the ability to view Veeam ONE dashboards and leverage reporting capabilities directly from within the Veeam Backup & Replication console. The VBR web UI now includes both a dashboards section and a reports section that surface Veeam ONE content, providing administrators with monitoring and reporting capabilities without leaving their primary backup management interface.

The dashboards section in the VBR web UI displays advanced dashboards for monitoring infrastructure and data protection operations. The reports section allows generating, viewing, and managing both predefined and custom reports available in Veeam ONE. This integration covers a wide range of reporting topics including backup inventory, repository usage, security compliance, capacity planning, change tracking, and operational trends.

This deeper integration reflects Veeam's strategy of making Veeam ONE an embedded component of the overall data protection workflow rather than a separate standalone tool. For administrators who primarily work within the VBR console, having monitoring dashboards and reporting immediately available eliminates context-switching and accelerates troubleshooting workflows.

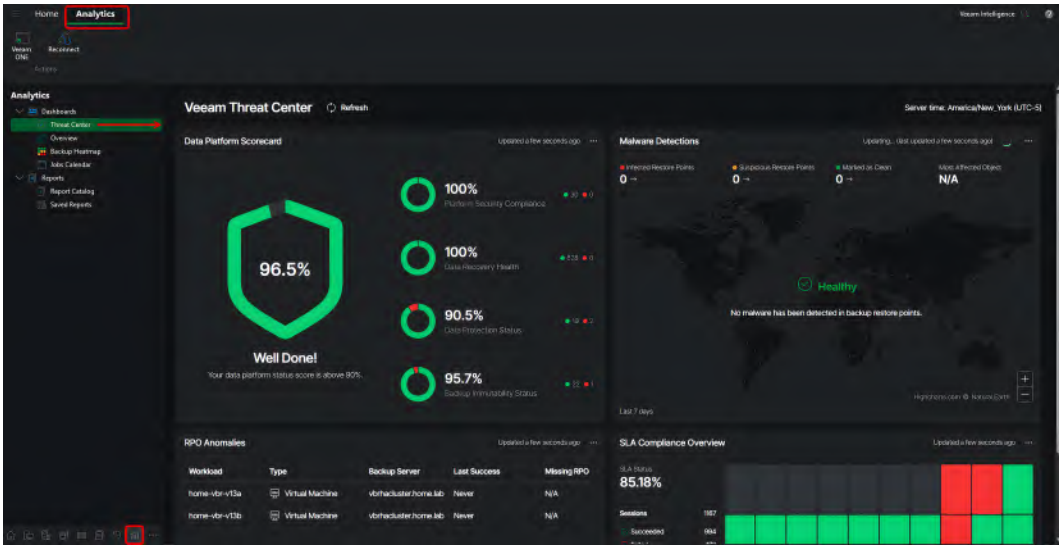


Figure 11.15: Veeam ONE integration in the VBR console

Key additions in version 13.0.1

The 13.0.1 update, released on November 19, 2025, brought several significant enhancements beyond the initial v13 release. These additions focus on expanded monitoring coverage, enhanced Veeam intelligence capabilities, and deeper platform integration.

Feature	Description
HA monitoring	Continuous health monitoring for VBR High-Availability clusters with failover alarms and automated remediation
Coveware recon in threat center	Recon findings with actionable alarms and filtering options for alarm fatigue reduction
Malware threat analysis agent	AI agent for real-time anomaly insights, blast radius assessment, and guided recovery
Deep data analysis agent	AI agent for conversational report generation and dataset navigation
Speech-to-Text input	Voice input for Veeam intelligence queries in the user's own language
Data resilience summary	Consolidated backup environment snapshot with aggregated status and recommendations
Enterprise app plug-in support	Expanded monitoring for SAP HANA, Oracle RMAN, SAP on Oracle, IBM Db2, and SAP MaxDB
Enhanced REST APIs	Optimized APIs with more cloud backup data and reduced execution time

Table 11.21: Key addition summary

The HA monitoring capability is particularly noteworthy. Veeam ONE 13.0.1 can continuously monitor VBR HA clusters and trigger automated remediation actions, including initiating a failover when the primary node is detected as unhealthy. This enables policy-driven failovers for both planned maintenance and unplanned outages, adding a new layer of operational resilience to the backup infrastructure.

Veeam ONE v13 and its 13.0.1 update represent a generational leap in backup infrastructure monitoring and observability. The combination of a redesigned web client, AI-powered Veeam

Technet24

intelligence agents, a PostgreSQL-backed interactive reporting engine, and deeper integration with both the VBR console and third-party platforms like ServiceNow creates a monitoring solution that is genuinely transformative for enterprise data protection operations.

For administrators planning their upgrade path, the transition to v13 is supported as an in-place upgrade from version 12.0.0.2498 and newer. The upgrade process introduces the new PostgreSQL reporting database as a required component, but the existing SQL server monitoring database is retained and automatically configured. The Veeam ONE agent is seamlessly transitioned to the new Veeam analytics service during the upgrade process.

The architectural direction of Veeam ONE v13 with its move toward PostgreSQL, HTML5-based reporting, and AI-powered analytics signals Veeam's long-term vision of a unified observability platform that goes beyond traditional backup monitoring. As the Veeam intelligence agents continue to evolve and the PostgreSQL migration eventually extends to all database functions, Veeam ONE is positioned to become the central intelligence hub for data protection operations across the enterprise.

Summary

This chapter has provided you with a detailed exploration of the new Veeam ONE v13 features and enhancements. We discussed what is new with Veeam ONE monitoring for v13 and looked at the Veeam ONE architecture and how it has changed from v12 to v13 with some new components and a new database. Lastly, we explored the new dashboard, looked into Veeam intelligence, and checked out the new reporting structure and reports available.

This chapter ensured that you understand everything that has changed or been enhanced in the new Veeam ONE v13, including the architecture changes. The next and final chapter of the book will cover some of the new recovery options available in Veeam Backup & Replication v13, including the new instant recovery to Azure.

Further reading

- Veeam ONE Documentation - `https://helpcenter.veeam.com/docs/one/userguide/about_one.html?ver=13`
- Sizing and Scalability Best Practices - `https://helpcenter.veeam.com/docs/one/userguide/scalability.html?ver=13`

12

Recovery Options

Veeam Backup & Replication v13 is part of the **Veeam Data Platform**, the newest release of Veeam and has many new recovery option enhancements. In this final chapter, we will discuss the new instant recovery to Azure and how to use it. We will look at **Universal CDP** and how this helps with recovery of your environment and servers. Lastly, we will look at the many other recovery options available to you, the backup administrator, for recovering your systems.

In this chapter, we're going to cover the following main topics:

- Instant recovery to Azure
- Universal CDP: Continuous data protection for near-zero recovery
- Veeam's other recovery options

Instant recovery to Microsoft Azure

One of the most significant new capabilities introduced in Veeam Backup & Replication v13 is **instant recovery to Microsoft Azure**. In previous versions of Veeam Backup & Replication, instant recovery was limited to on-premises hypervisor platforms including VMware vSphere, Microsoft Hyper-V, and Nutanix AHV. Version 13 breaks this boundary by extending the instant recovery framework directly into the public cloud, enabling administrators to boot workloads as native Azure virtual machines within minutes, regardless of whether the original source was a VMware VM, a Hyper-V VM, an Amazon EC2 instance, a Nutanix AHV VM, or a physical server protected by Veeam agent.

This section explores what instant recovery to Azure is, how the underlying architecture works, the prerequisites and considerations you need to plan for, the step-by-step recovery wizard workflow, and the finalization process that transitions a temporarily mounted workload into a fully production-ready Azure VM.

Technet24

What is instant recovery to Azure?

Instant recovery to Microsoft Azure is a restore operation that allows you to immediately recover a wide range of Veeam-protected workloads as native Azure VMs. During recovery, Veeam Backup & Replication does not extract the full workload data from the backup file to Azure managed disks. Instead, it creates a "dummy" VM in Azure and mounts the workload's virtual disks directly from the backup stored in Azure-based object storage. Because there is no need to fully hydrate the VM's disks before booting, the workload can be online in a matter of minutes, often under five minutes, including Azure VM provisioning time and the boot sequence to the operating system logon screen.

This approach draws on the same fundamental instant recovery concepts Veeam has offered for on-premises hypervisors since the introduction of vPower NFS technology in earlier versions. The key difference is that the mount target is now an Azure virtual machine rather than an ESXi or Hyper-V host, and the backup data source must reside in an Azure-based repository. The dummy VM will have limited I/O performance because all disk reads are serviced from the compressed and deduplicated backup file over the network. However, for disaster recovery failover, DR testing, or temporary access to application data, this trade-off is highly advantageous compared to waiting hours for a full restore.

> **Note**
>
> **Key concept:**
>
> Instant recovery to Azure is not a replacement for a full restore. It is designed to minimize RTO by providing rapid, temporary access to workloads. For sustained production use, you must finalize the recovery by migrating the VM to production (switchover), which copies the data to native Azure managed disks and removes the helper appliance dependency.

What's new in v13: From on-premises only to Azure-native recovery

Prior to version 13, instant recovery in Veeam Backup & Replication was constrained to three on-premises hypervisor targets: **VMware vSphere** (using vPower NFS), **Microsoft Hyper-V**, and **Nutanix AHV**. If an administrator needed to failover workloads to the public cloud, the only option was the full "*Restore to Microsoft Azure*" operation, which required extracting the entire backup, converting disks, and provisioning the VM—a process that could take considerable time depending on backup size and available bandwidth. Version 13 changes this paradigm by introducing Azure as a first-class instant recovery target.

The key innovations include the following:

- **Azure as an instant recovery target**: Any image-level backup of a Windows or Linux workload can now be instantly recovered as a native Azure VM. This includes VMware VMs, Hyper-V VMs, Nutanix AHV VMs, Amazon EC2 instances, Proxmox VE VMs, and physical servers protected by Veeam agent.
- **Helper appliance template architecture**: A new helper appliance template system publishes pre-built Windows and Linux appliance images to the Azure compute gallery, significantly accelerating recovery start times by eliminating the need to deploy a fresh helper VM from scratch for each operation.
- **Sub-five-minute recovery times**: Veeam has benchmarked instant recovery to Azure at under five minutes from initiation to OS logon screen, inclusive of Azure VM provisioning and boot time.
- **Bulk recovery support**: Instant recovery to Azure supports recovering multiple workloads simultaneously, with up to 150 concurrent restore sessions. For larger-scale recovery, workloads can be processed in batches.
- **PowerShell automation**: New Veeam PowerShell cmdlets for instant recovery to Azure enable fully scripted and automated recovery operations, enabling integration with DR orchestration workflows.
- **Broader use cases beyond DR**: Because instantly recovered VMs run as native Azure VMs, the feature provides a foundation for clean room recovery, malware scanning, data classification, and DR testing in an isolated Azure environment.

Architecture and how it works

Understanding the architecture behind instant recovery to Azure is essential for both planning and troubleshooting. The process involves several components working together: the Veeam Backup & Replication server, an Azure-based backup repository, helper appliance templates stored in the Azure compute gallery, and temporarily deployed helper appliance VMs in Azure.

Helper appliance templates

Before you can perform your first instant recovery to Azure, you must create helper appliance templates. These templates are pre-configured VM images—one Windows-based and one Linux-based—that are published to the Azure compute gallery in your selected Azure subscription and region. The template creation process is a one-time setup step performed from the Veeam Backup & Replication console under **Backup Infrastructure** | **Helper Appliance Templates**.

During template creation, Veeam Backup & Replication deploys two temporary VMs in Azure (one Windows, one Linux) using the `Standard_D4ls_v5` VM size. It uploads the required

Technet24

Veeam components and services to these VMs, captures generalized images from them, publishes the images to the Azure compute gallery, and then removes the temporary VMs. Only the gallery images remain, which are cost-effective to store and serve as rapid deployment sources for future recovery operations. This process typically takes approximately 30 minutes to complete.

The recovery process flow

When you initiate an instant recovery to Azure, the following sequence of events occurs:

1. Veeam Backup & Replication deploys a helper appliance VM in Azure from the pre-built template in the Azure compute gallery. A Windows helper appliance is used when restoring Windows workloads, and a Linux helper appliance is used for Linux workloads.
2. The helper appliance connects to the Azure-based backup repository and reads the backup data for the selected restore point.
3. A new "dummy" Azure VM is created with the specified configuration (VM size, resource group, virtual network, subnet, and network security group)
4. The workload's virtual disks are mounted from the backup file to the dummy VM via the helper appliance, which acts as an intermediary that translates backup data into disk-level reads.
5. The dummy VM boots and becomes accessible as a native Azure VM. Any writes or changes made while the VM is running are tracked in redo logs managed by the helper appliance.
6. The helper appliance remains active for the duration of the instant recovery session, serving disk I/O requests and tracking changes. It is automatically removed only after the recovery is finalized.

Supported backup sources and repository types

Instant recovery to Azure supports a broad range of Veeam-protected workloads. However, the backups must reside in an Azure-based repository for the operation to work. The following table summarizes the supported workload types and repository types.

Workload type	Source platform
VMware vSphere VMs	Backups created by VBR
Microsoft Hyper-V VMs	Backups created by VBR
Nutanix AHV VMs	Backups via Veeam plug-in for Nutanix AHV

Workload type	Source platform
Proxmox VE VMs	Backups created by VBR
Amazon EC2 instances	Backups via Veeam backup for AWS
Physical servers (Windows)	Veeam agent for Microsoft Windows
Physical servers (Linux)	Veeam agent for Linux
VMware cloud director VMs	Backups created by VBR

Table 12.1: Supported workload types

Along with the supported workload types in the table, *Table 12.2* lists the supported repository types.

Repository type	Notes
Azure Blob storage (object storage repository)	Added as a backup repository in VBR
Veeam data cloud vault	Object storage added as backup repository
External repository (Veeam backup for Azure)	Backups created by Veeam backup for Azure

Table 12.2: Supported Azure Repository Types

Note

Important:

Backups stored on-premises (on local repositories, NFS/SMB shares, or deduplication appliances) cannot be used for instant recovery to Azure. The backup data must already reside in an Azure-accessible repository. If your backups are on-premises, you must first create backup copies to an Azure Blob storage repository or Veeam data cloud vault before instant recovery to Azure becomes available.

Prerequisites and considerations

Before performing instant recovery to Azure, ensure your environment meets the following requirements:

Technet24

Azure account and subscription requirements

You should have the following access:

- An Azure compute account must be registered in Veeam Backup & Replication with the appropriate permissions to create VMs, managed disks, network interfaces, and Azure compute gallery resources.
- The Azure subscription must have sufficient CPU quotas available for both the helper appliance VMs (which use `Standard_D4ds_v5` by default, consuming 4 vCPUs from the Dldsv5 series) and the recovered workload VMs.
- Storage accounts used for helper appliance template creation and communication must have key-based authentication enabled. Storage accounts with private endpoints are not supported.
- Sufficient available IP addresses in both the production subnet (for the recovered VM) and the appliance subnet (for the helper appliance) are required.

Operating system and workload constraints

You should keep the following constraints in mind:

- Windows workloads must run Windows server 2016, Windows 10, or later. Windows workloads configured with the BIOS boot type that have Windows updates from March 2025 or newer will not boot after instant recovery to Azure.
- Linux workloads must use a distribution supported by Microsoft Azure (refer to the Azure supported distributions documentation). The workload must have dracut, mkinitrd, or initramfs installed, and the iSCSI initiator package must be present (open-iscsi for Debian/Ubuntu/SUSE, iscsi-initiator-utils for RHEL-based).
- Linux workloads with system mount points residing on ZFS pools are not supported.
- Windows workloads where the boot partition and operating system partition reside on separate drives are not supported.
- BitLocker-encrypted VMs are not supported for restore from the Windows-based backup server (except when restoring from Veeam agent for Windows backups).
- Instant recovery of the VSA itself is not supported.
- For optimal performance, avoid keeping multiple Linux kernel versions installed on workloads you plan to recover. During recovery, Veeam rebuilds initramfs/dracut for each detected kernel, and multiple kernels will significantly increase conversion time.

Network considerations

If you select an existing network and subnet for the helper appliance or helper appliance template, you must allow access to public endpoints from those subnets. This is critical due to the new Microsoft Azure changes: after March 31, 2026, new virtual networks default to private subnets that no longer support default outbound access connectivity. VMs requiring public endpoint access will need explicit outbound connectivity configured. If you use the default option to create a temporary virtual network and subnet for the helper appliance template, this restriction does not apply. In this case, the created temporary subnet will retain allowed outbound internet access. It is also highly recommended to configure a separate subnet for appliance traffic, as some workloads may fail to boot if the appliance and production traffic share the same subnet.

Instant recovery to Azure wizard

Instant Recovery to Azure is initiated from the Veeam Backup & Replication console.

> **Note**
>
> Windows console in v13 and v13.0.1; the Web UI does not yet support this operation.

To begin, follow these steps:

1. Navigate to the **Home** view, click the **Restore** button in the toolbar, and select **VMware vSphere | Restore from backup | Entire VM restore | Instant recovery | Instant recovery to Microsoft Azure**.



Figure 12.1: Instant recovery to Microsoft Azure option

Technet24

The following wizard steps guide you through the configuration:

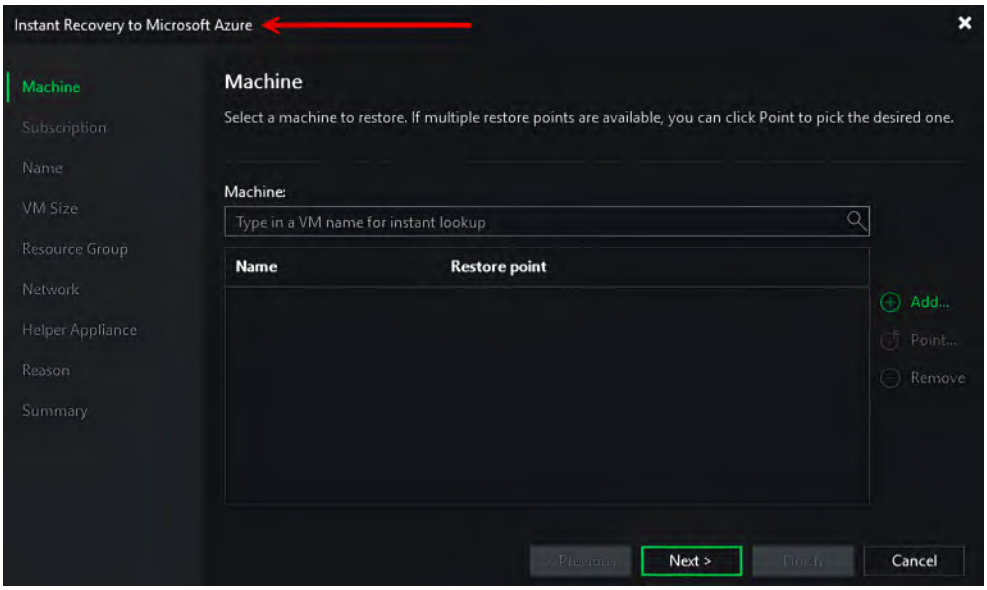


Figure 12.2: Instant recovery to Microsoft Azure wizard

The steps shown in *Figure 12.2* are explained in *Table 12.3*. (I do not have an Azure subscription, so I cannot proceed through the wizard for this section.)

Wizard step	Configuration details
Restore point	By default, the most recent available restore point is selected. Click **Point...** to choose a different restore point from the available backup chain.
Subscription and region	Select the Azure subscription and Azure region where the recovered VM will be created.
VM configuration	Configure the VM name (can be modified from the source name), add Azure tags, select the VM size (defaults to matching the source vCPU/RAM), and choose the disk type. Note that ultra disks are not supported; available options include standard HDD, standard SSD, and premium SSD.

Wizard step	Configuration details
Resource group	Select an existing resource group or allow Veeam to create a new one. For mass recovery scenarios, selecting an existing resource group is recommended to reduce Azure API resource creation overhead.
Network	Choose the virtual network, subnet, network security group (NSG), and whether to assign a public IP address to the recovered VM.
Helper appliance	Select the subnet where the helper appliance will be deployed and the storage account it will use. For parallel restores, select an existing storage account rather than allowing Veeam to create a temporary one, as Azure limits storage account creation rates.
Summary	Review all selected options. Optionally enable verification to test whether the recovered VM boots successfully by attempting an RDP or SSH connection. Ensure the associated NSG rules allow RDP/SSH if this option is selected.

Table 12.3: Instant recovery to Azure wizard steps

2. Click **Finish**. Veeam Backup & Replication deploys the helper appliance from the compute gallery template, creates the target VM in Azure, and mounts the workload's disks from the backup. This process is fast regardless of the size of the workload because no actual data restore takes place—only disk mount operations are performed against the backup data in the Azure repository.

Finalizing instant recovery: Migrate or discard

An instant recovery session cannot remain active indefinitely. The helper appliance continues to consume Azure compute resources and serves all disk I/O for the recovered VM. You must finalize the recovery by choosing one of two paths:

- **Migrate to production (switchover)**: This option copies the workload data from the backup to native Azure managed disks, applies any changes (redo logs) that accumulated while the VM was running in instant recovery mode, and then removes the helper appliance. After switchover, the VM becomes a fully independent,

Technet24

production-ready Azure VM with full I/O performance. The switchover can be triggered immediately or scheduled for a specific date and time.

- **Stop publishing (discard)**: This option removes the recovered VM and the helper appliance from Azure entirely. All changes made during the instant recovery session are lost. This path is ideal for DR testing scenarios where you simply need to verify that the workload boots correctly and applications are functional in Azure, without intending to keep the VM running.

> **Note**
>
> **Best practice:**
>
> For disaster recovery testing, use the discard option to validate recovery without incurring ongoing Azure VM costs. For actual failover scenarios, always complete the switchover promptly to transition the workload to full-performance Azure managed disks and eliminate the helper appliance dependency.

Mass recovery considerations

For large-scale disaster recovery scenarios involving dozens or hundreds of workloads, instant recovery to Azure supports bulk processing with up to 150 concurrent restore sessions. When planning mass recovery, several additional considerations apply to ensure reliability and avoid Azure API throttling or resource exhaustion.

Planning area	Recommendation
Concurrent limit	Maximum of 150 simultaneous restores. Process additional workloads in batches after completing switchover for the first batch.
Resource groups	Select existing resource groups to reduce the number of Azure resources created during each restore operation.
Storage accounts	Select an existing storage account for the helper appliance rather than creating temporary accounts. Azure limits the rate of storage account creation, which can cause failures in parallel restores.

Planning area	Recommendation
Network security groups	Select existing NSGs rather than using the default empty option, which creates a new NSG per workload.
Public IP addresses	Do not assign public IPs during mass recovery. Each public IP is an additional Azure resource that increases API load.
Appliance subnet	Use a dedicated, separate subnet for helper appliance traffic. Sharing a subnet with production workloads can cause boot failures.
CPU quotas	Verify that the Azure subscription has sufficient vCPU quotas for both the recovered VMs and the helper appliance VMs (4 vCPUs each from the Dldsv5 series by default).

Table 12.4: Mass recovery planning checklist

Instant recovery to Microsoft Azure in Veeam Backup & Replication v13 represents a fundamental expansion of Veeam's recovery capabilities. By extending the proven instant recovery framework from on-premises hypervisors to the Azure public cloud, Veeam v13 enables organizations to achieve sub-five-minute RTOs for cloud-based disaster recovery, eliminate the lengthy data hydration process required by traditional full restores to Azure, leverage Azure as a DR testing sandbox without permanent resource commitment, scale recovery operations to 150 concurrent workloads for enterprise-grade failover scenarios, and automate recovery workflows using new PowerShell cmdlets purpose-built for azure instant recovery.

The helper appliance template architecture and Azure Compute Gallery integration ensure that recovery start times remain consistently fast, while the switchover finalization process provides a clear path from temporary failover to permanent production operation. Combined with the broader v13 platform enhancements including the Veeam Software Appliance, AI-powered malware detection, Universal CDP, and the new Web UI, instant recovery to Azure positions Veeam Backup & Replication as a comprehensive cyber-resilience platform capable of meeting the recovery demands of modern hybrid-cloud environments.

Universal CDP: Continuous data protection for near-zero recovery

Continuous data protection (CDP) has long been the gold standard for minimizing data loss in mission-critical environments. With Veeam Backup & Replication v13, the CDP engine has

Technet24

been completely rebuilt as **Universal CDP**—a hypervisor-agnostic, scalable architecture that decouples real-time replication from any single virtualization platform. This section examines how Universal CDP works under the hood, what has changed from the vSphere-only CDP of earlier versions, and how administrators can leverage it to achieve recovery point objectives (RPOs) measured in seconds rather than hours.

From vSphere CDP to Universal CDP

In Veeam Backup & Replication v12 and earlier, CDP was tightly coupled to VMware vSphere through the vSphere I/O filter framework (VAIO). While effective, this created significant constraints. CDP was limited exclusively to VMware environments, required VAIO filter installation on every ESXi host, and introduced dependencies on vCenter availability for filter lifecycle management. Organizations running mixed hypervisor estates or those migrating toward platform diversity had no way to extend CDP protection beyond vSphere boundaries.

Version 13 introduces a fundamentally different approach. Universal CDP replaces the VAIO-dependent architecture with a **platform-independent I/O journal engine** that operates at the Veeam proxy layer. Instead of relying on hypervisor-native I/O filter frameworks, Universal CDP intercepts and journals write operations through Veeam's own data movers. This architectural shift delivers several immediate advantages:

- **Hypervisor independence**: CDP policies can now protect workloads running on VMware vSphere, Microsoft Hyper-V, Nutanix AHV, and other supported platforms using a single, unified engine.
- **Eliminated VAIO dependency**: No I/O filter installation or lifecycle management is required on hypervisor hosts, reducing operational overhead and removing a potential failure domain from the data path.
- **Simplified proxy-based deployment**: CDP data processing is consolidated at the proxy tier, which aligns with Veeam's broader v13 architecture of centralizing intelligence in the proxy and repository layers.
- **Scalable journal management**: The I/O journal engine has been redesigned for horizontal scalability, supporting higher transaction volumes and larger journal retention windows without proportional increases in resource consumption.

Note

Key distinction:

Universal CDP in v13 is not merely a rebranding of the previous CDP feature. The underlying I/O interception, journaling, and replay mechanisms have been rebuilt from

> the ground up. Existing CDP policies from v12 must be re-created after upgrading, as the data path architecture is entirely new.

Universal CDP architecture and components

Understanding how Universal CDP components interact is essential for proper sizing, placement, and troubleshooting. The architecture consists of four primary layers, each responsible for a distinct phase of the continuous protection pipeline.

I/O journal engine

The **I/O journal engine** is the core of Universal CDP. It maintains a chronologically ordered, write-ahead log of every write operation issued to protected virtual disks. Each journal entry records the exact byte offset, payload, and timestamp of the write, enabling point-in-time recovery to any moment within the journal retention window.

In v13, the journal engine has been optimized for high-throughput workloads with several key improvements:

- **Parallel journal streams**: Multiple virtual disks on the same VM can journal concurrently without serialization, eliminating a bottleneck present in earlier implementations.
- **Compressed journal entries**: Write payloads are compressed inline before being written to the journal store, reducing storage consumption by 40–60% for typical enterprise workloads.
- **Journal checkpointing**: Periodic consistency checkpoints allow the journal to be truncated without replaying the entire history, accelerating both recovery operations and journal maintenance.
- **Configurable retention granularity**: Administrators can define journal retention in seconds, minutes, or hours, with separate short-term and long-term retention tiers that automatically coalesce older journal entries to reduce storage footprint.

CDP proxy

The **CDP proxy** serves as the data mover for continuous protection. It intercepts I/O from protected workloads, feeds write operations into the journal engine, and manages the replication stream to the target. In v13, CDP proxies are standard Veeam Linux proxies and there is no separate CDP proxy role. Any Linux proxy with sufficient resources can be designated for CDP duty, which simplifies infrastructure planning and allows organizations to repurpose existing proxy capacity.

Technet24

For optimal performance, CDP proxies should be placed as close as possible to the protected workloads in terms of network latency. Each proxy can handle multiple concurrent CDP policies, but administrators should monitor CPU and memory utilization carefully, as the journaling process requires both compute resources for compression and memory for the in-flight write buffer.

CDP target and replica datastores

Universal CDP replicates protected VMs to a target datastore, maintaining both a base replica image and the I/O journal. The target can reside on any supported storage accessible to the target hypervisor, including local datastores, SAN/NAS storage, or deduplication appliances. The journal files are stored alongside the replica on the target, occupying additional space proportional to the write change rate and the configured retention window.

Component	Sizing consideration	v13 recommendation
CDP proxy CPU	Journal compression and I/O processing	4 vCPUs per 10 concurrently protected VMs
CDP proxy memory	In-flight write buffer and journal indexing	8 GB base + 512 MB per protected VM
Journal storage	Change rate × retention window	2× daily change rate × retention (hours)
Network bandwidth	Sustained write throughput to target	10 GbE minimum; 25 GbE for >50 VMs
Target datastore	Base replica + journal files	VM provisioned size + journal estimate

Table 12.4: CDP component sizing guidance

Network resilience and I/O journal continuity

One of the most significant improvements in Universal CDP is its handling of network disruptions between the source and target sites. In previous versions, a network outage during CDP replication could force a full resynchronization of the replica, effectively negating the RPO benefits of continuous protection during the recovery window.

Version 13 introduces a **local journal spooling mechanism** on the source-side CDP proxy. When network connectivity to the target is lost, the proxy continues capturing and journaling write operations locally. The local journal spool acts as a buffer, preserving the complete I/O

history for the duration of the outage. Once connectivity is restored, the proxy replays the buffered journal entries to the target in chronological order, bringing the replica back to a fully current state without requiring a resynchronization of the base image.

This behavior is governed by two configurable parameters:

- **Local spool size limit**: Defines the maximum amount of storage the proxy will consume for local journal buffering. If the spool reaches this limit before connectivity is restored, the CDP policy enters a degraded state and will require resynchronization. Default is 10 GB.
- **Outage tolerance timer**: Specifies the maximum duration (in minutes) the proxy will continue local spooling before automatically triggering a resynchronization. This prevents unbounded spool growth for extended outages. Default is 60 minutes.

Tip

Size your local spool based on your environment's peak write change rate multiplied by the maximum tolerable outage duration. For a workload generating 500 MB/hour of change data, a 60-minute outage tolerance requires at least 500 MB of spool space, plus a 20% safety margin.

The network resilience improvements extend to WAN-based CDP scenarios as well. Universal CDP supports configurable WAN acceleration profiles that optimize journal replication across high-latency, bandwidth-constrained links. Journal entries are deduplicated and compressed before transmission, and the replication stream uses an adaptive windowing protocol that adjusts transfer rates based on real-time network conditions.

Recovery operations with Universal CDP

The primary value proposition of CDP is the ability to recover to a precise point in time with minimal data loss. Universal CDP in v13 provides multiple recovery methods, each suited to different operational scenarios. Understanding when and how to use each method is critical for meeting recovery time objectives (RTOs) alongside the near-zero RPOs that CDP delivers.

Point-in-time failover

Point-in-time failover is the most common recovery operation for CDP-protected workloads. When a failure occurs on the source VM, the administrator initiates a failover to the CDP replica, selecting a specific point in time from the I/O journal. The granularity of available restore points depends on the write activity of the source VM. Every write operation creates a potential recovery point, meaning that for active workloads, recovery points may be available at sub-second intervals.

Technet24

The failover process in v13 follows a streamlined workflow:

1. Select the CDP replica from the **Home** view under **Replicas | Ready**. Right-click and choose **Failover Now**.
2. Choose the restore point using the timeline slider or by entering a precise timestamp. The console displays the RPO achievable at each point, accounting for any journal gaps from network interruptions.
3. Journal replay begins automatically. The CDP engine applies all journal entries from the last checkpoint up to the selected point in time, reconstructing the VM's disk state at that exact moment.
4. The replica VM powers on at the target site with the recovered state. Network re-IP rules, if configured in the CDP policy, are applied automatically to adapt the VM to the target site's network configuration.
5. Validate and commit. After verifying the recovered VM is operational, the administrator commits the failover, which discards the pre-failover state, or undoes it to revert to the original replica state.

Planned failover

Planned failover is used for scheduled migrations or maintenance windows where the source VM can be gracefully shut down before the failover. In this scenario, the CDP engine ensures that all pending journal entries are fully replicated to the target before the source VM is powered off, guaranteeing zero data loss. Planned failover also performs automatic verification of the target replica's consistency before powering on the recovered VM, providing an additional safety check that is not available during an unplanned failover.

Failback and reverse protection

After a failover event, the recovered VM is running at the target site on the CDP replica. To return to the original production site, v13 provides integrated failback with **reverse CDP protection**. Once the source site is restored, the administrator initiates failback, which performs the following actions:

- Synchronize changes made on the target-side VM back to the original source location using the same CDP journal mechanism.
- Optionally performs a quick rollback if the source VM's disks are still intact, replaying only the delta changes rather than transferring the full VM image.
- Re-establishes CDP protection in the original direction automatically after failback completes, restoring the environment to its pre-failure protection posture without manual policy recreation.

Balancing RPO against infrastructure overhead

While Universal CDP delivers the lowest possible RPO of any Veeam protection mechanism, it does so at a cost in infrastructure resources. Every write operation to a protected VM is intercepted, compressed, journaled, and replicated in near-real time. For high-transaction workloads such as database servers or email systems, this can generate substantial and sustained load on CDP proxies and network links.

Veeam v13 provides several controls to help administrators balance protection granularity against resource consumption:

RPO tier	Journal behavior	Typical use case	Resource impact
Sub-second	Every write journaled immediately	Tier-0 databases, financial systems	Highest – dedicated proxy, 25 GbE
Seconds (5–30s)	Micro-batch journal writes	Application servers, ERP systems	High – shared proxy feasible
Minutes (1–5m)	Periodic journal checkpoints	File servers, development environments	Moderate – standard proxy capacity
Short-term only	CDP for hours, then traditional backup	Compliance-driven short RPO windows	Lower – limited journal retention

Table 12.5: RPO tiers and resource impact

The micro-batch journaling option is new to v13 and represents a pragmatic middle ground. Rather than journaling every individual write synchronously, the CDP engine batches writes over a short configurable interval (5–30 seconds) and journals them as a single atomic entry. This reduces the I/O overhead on both the proxy and the network by 30–50% while still delivering RPOs far superior to traditional snapshot-based backup. For many enterprise workloads, the few seconds of additional potential data loss are an acceptable trade-off for significantly reduced infrastructure requirements.

Universal CDP in the context of the broader recovery strategy

It is important to understand that Universal CDP is not a replacement for traditional backup. CDP and backup serve complementary roles in a comprehensive data protection strategy, and v13 is explicitly designed to operate both mechanisms in parallel.

Technet24

CDP protects against operational failures such as hardware crashes, application corruption, and accidental deletions where the goal is to recover to the most recent possible state with minimal data loss. The CDP replica exists as a live, continuously updated copy of the source VM, making it ideal for rapid failover scenarios. However, CDP replicas are not immutable. If ransomware encrypts the source VM's disks, those encrypted writes will be faithfully replicated to the CDP target. While the I/O journal allows recovery to a point in time before the encryption began, the journal retention window is finite, and a sophisticated attack that lies dormant beyond the retention window could compromise the replica.

Traditional backup, by contrast, creates isolated, point-in-time copies that can be stored on immutable repositories, air-gapped media, or offsite object storage with extended retention. These backups serve as the last line of defense against scenarios that CDP alone cannot address, such as extended ransomware dwell times, insider threats, compliance-driven long-term retention requirements, and disaster recovery to geographically distant sites.

Note

Best practice:

Always pair Universal CDP with regular backup jobs targeting an immutable or hardened repository. Use CDP for operational recovery with the shortest possible RPO, and use backup for strategic recovery, long-term retention, and ransomware resilience. This dual-layer approach is the foundation of Veeam's recommended data protection strategy in v13.

Table 12.6 compares Universal CDP with traditional backup:

Characteristic	Universal CDP	Traditional backup
Recovery point objective	Seconds to minutes	Hours (last scheduled backup)
Recovery time objective	Minutes (VM failover)	Minutes to hours (restore)
Protection mechanism	Continuous I/O journaling	Periodic snapshot + transfer
Ransomware resilience	Limited to journal retention window	Strong with immutable repositories

Characteristic	Universal CDP	Traditional backup
Long-term retention	Not designed for this purpose	Days, months, years with GFS
Infrastructure cost	Higher (always-on replication)	Lower (scheduled processing)
Platform support (v13)	All supported hypervisors	All supported hypervisors

Table 12.6: Universal CDP vs. traditional backup comparison

Monitoring and troubleshooting CDP health

Universal CDP introduces dedicated monitoring surfaces in both the Veeam Backup & Replication console and Veeam ONE. Effective monitoring is essential because CDP is a continuous process, unlike scheduled backup jobs, which run and complete. CDP policies must operate uninterrupted to maintain their protection guarantees.

Key metrics to monitor

The following metrics should be tracked for every CDP policy in a v13 environment:

- **Current RPO**: The actual time gap between the last successfully replicated journal entry and the present moment. This metric reveals whether the CDP policy is meeting its RPO target in real time.
- **Journal lag**: The volume of journal data awaiting replication to the target. A growing lag indicates that the network or target storage is unable to keep pace with the source's write rate.
- **Local spool utilization**: Tracks consumption of the source-side journal spool during network outages. Approaching the spool limit is an early warning that a resynchronization may be imminent.
- **Proxy resource utilization**: CPU, memory, and network throughput on CDP proxies. Sustained high utilization suggests the proxy is undersized for its workload.
- **Journal storage consumption**: Total disk space consumed by journal files on the target datastore, tracked against available capacity to prevent journal exhaustion.

Common troubleshooting scenarios

CDP is a great tool for recovery, but what happens when things do not work as expected or you encounter errors? The following are some common troubleshooting scenarios:

- **RPO violation alerts**: If the reported RPO exceeds the configured target, check journal lag first. If lag is growing, the bottleneck is typically network bandwidth or target

Technet24

storage write performance. If lag is stable but RPO is elevated, the issue may be proxy resource contention. In such cases, verify that the CDP proxy is not also handling heavy backup or restore traffic simultaneously.

- **Journal spool exhaustion**: When the local spool fills during a network outage, the CDP policy cannot continue journaling and will require a full resynchronization when connectivity returns. To prevent this, increase the spool size limit in the CDP policy settings and ensure the proxy has sufficient local storage provisioned for spool capacity.
- **Replica consistency check failures**: During planned failover, the consistency check may report issues if the target datastore experienced write errors or if journal entries were corrupted due to storage-level issues. In such cases, review the target storage health, check for any SMART errors or LUN latency spikes, and consider running a manual resynchronization to rebuild the replica from a known-good state.

Universal CDP in Veeam Backup & Replication v13 represents a generational leap in continuous data protection. By decoupling the CDP engine from hypervisor-specific I/O filter frameworks, Veeam has created a platform-independent protection mechanism that can deliver near-zero RPOs across diverse infrastructure environments. The redesigned I/O journal engine, network resilience capabilities, and flexible recovery options give administrators the tools they need to protect their most critical workloads against data loss with unprecedented granularity.

However, CDP is most effective when deployed as part of a layered protection strategy. Pairing Universal CDP with traditional backup to immutable or hardened repositories ensures that organizations are protected not only against operational failures requiring rapid recovery, but also against sophisticated threats like ransomware and insider attacks that require the safety net of isolated, long-retention backup copies.

Veeam's other recovery options

While instant VM recovery and full VM restore are the most commonly used recovery paths in Veeam Backup & Replication v13, the platform provides a rich set of additional recovery capabilities designed to address granular, targeted, and specialized scenarios. Whether you need to recover a single deleted file, restore a corrupted volume, extract an individual email from an exchange backup, or bring an entire physical machine back from bare metal, Veeam v13 has a purpose-built recovery option to match.

In this section, we will cover the following recovery types beyond instant VM recovery and full VM restore:

- **File-level recovery** (**FLR**): Restore individual files and folders from image-level backups.

- **Volume-level restore**: Restore one or more volumes to their original or alternate location.
- **Virtual disk restore**: Restore specific VM virtual disks without restoring the entire VM.
- **Application item restore** (**Veeam explorers**): Granular recovery for Microsoft AD, SQL Server, Exchange, SharePoint, Oracle, and PostgreSQL.
- **Export backup**: Synthesize a portable, self-contained full backup from any restore point chain.
- **Secure restore**: Scan backup contents for malware before data is returned to production.
- **Instant cloud recovery to Microsoft Azure**: A v13 feature enabling rapid recovery of any workload as a native Azure VM.
- **Bare metal recovery** (**BMR**): Full physical or virtual machine recovery using Veeam recovery media.

File-level recovery (FLR)

File-level recovery (**FLR**) is one of the most frequently used restore paths in daily operations. It allows administrators or delegated restore operators to retrieve individual files and folders from an image-level backup without restoring the entire virtual machine or physical machine. The source VM can remain fully online and operational during the process.

Veeam v13 introduces Linux OS mount server support for FLR, meaning a dedicated Linux mount server can now be used to mount and browse backup contents for both Windows and Linux guest operating systems. This completes the transition to a fully Linux-capable backup infrastructure and removes the previous dependency on a Windows mount server.

Technet24

How FLR works

When FLR is initiated, Veeam mounts the disks of the backed-up machine to either the backup server, a designated mount server, or a helper host (for Linux guest VMs). The Veeam backup browser then opens and presents the full directory structure of the backed-up machine at the selected restore point. The operator browses the tree, selects the required files or folders, and restores them to the original location, an alternate path, or downloads them directly.

The FLR process follows these steps:

1. Right-click the backed-up machine in **Home** | **Backups** and select **Restore guest files...**

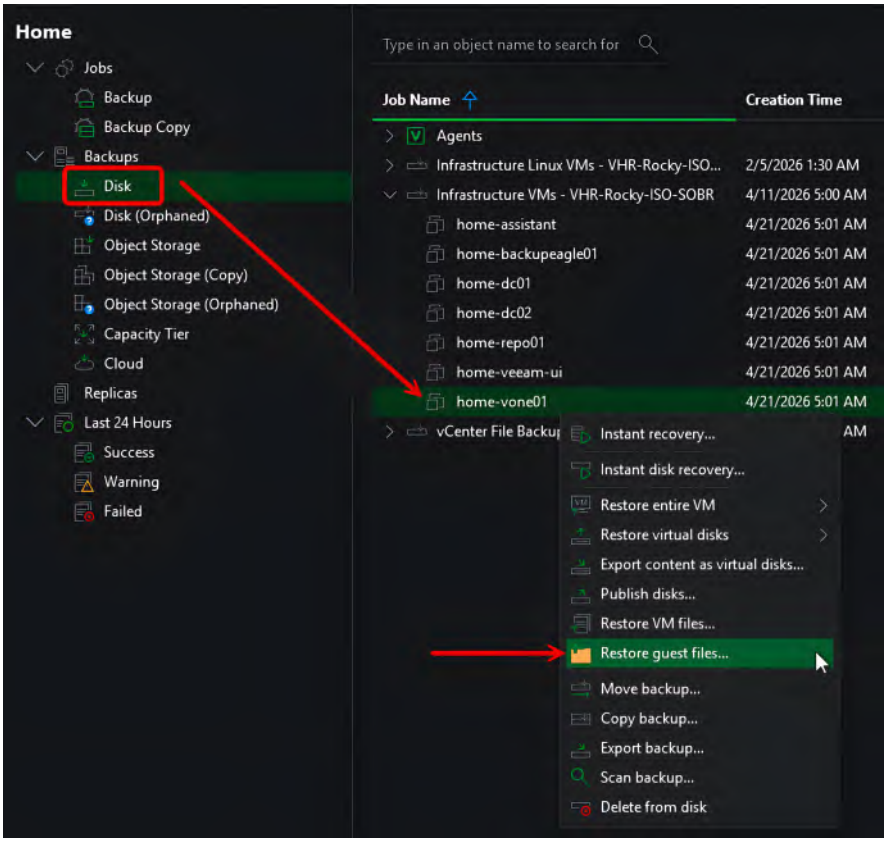


Figure 12.3: Right-click VM to show restore guest files

2. Select the restore point to browse.

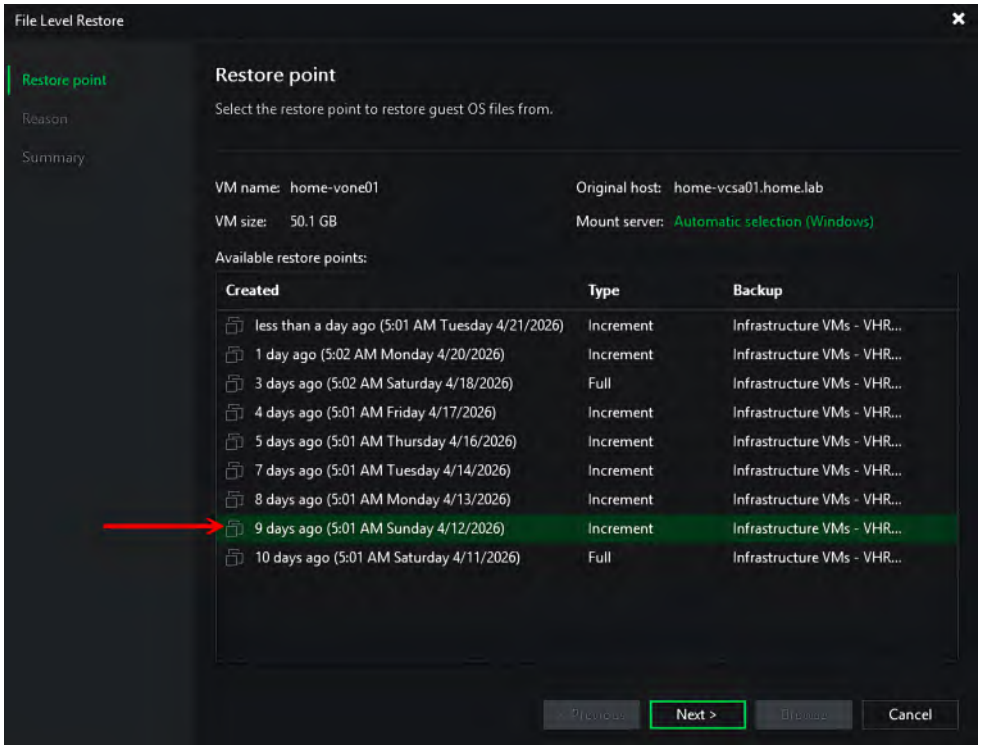


Figure 12.4: Selecting the restore point to restore from

The Veeam backup browser mounts the backup and presents the guest file system.

3. Navigate to the required files or folders, right-click, and choose **Restore to** or **Copy to**.

Technet24

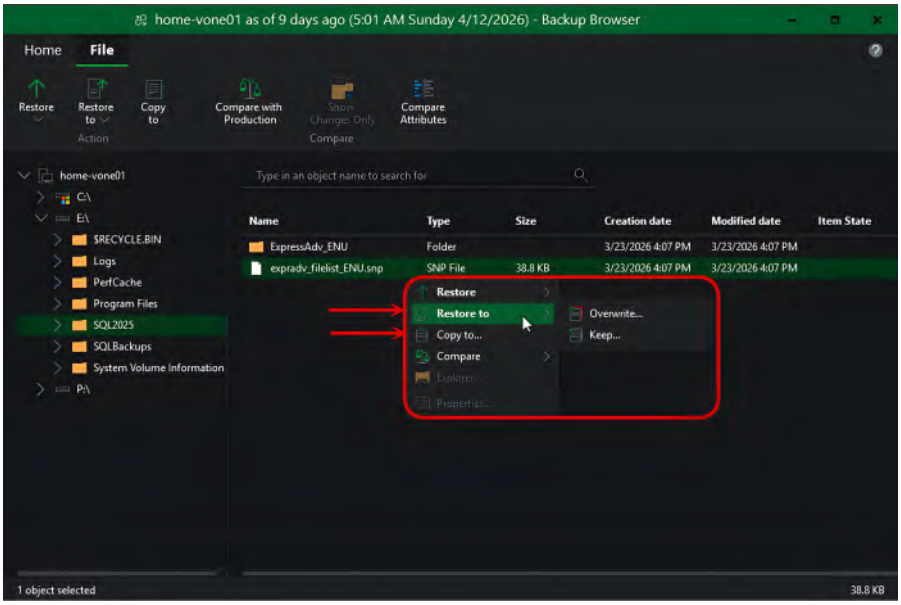


Figure 12.5: Selecting Restore or Copy to option

4. Specify the destination (original location, custom path, or download)

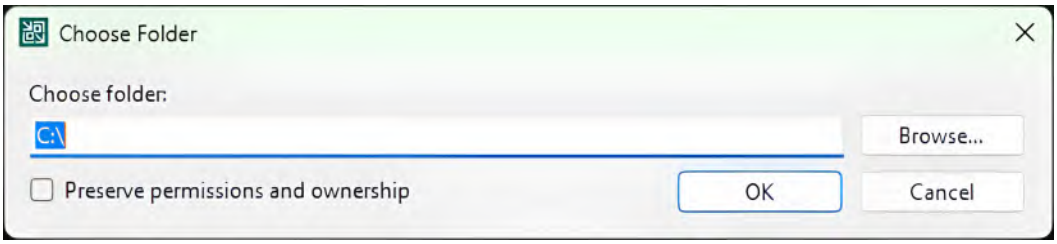


Figure 12.6: Specify destination of restore/copy

Veeam transfers the selected data and dismounts the backup when complete or after a session timeout.

Windows vs. Linux FLR

For Windows guest VMs, Veeam mounts the backup content using its proprietary vPower NFS driver or a direct mount to the mount server. Administrators can restore files with or without network connectivity to the guest OS, which is a networkless restore path using the VMware Tools (VIX) channel when the network is unavailable. For Linux guest VMs, Veeam deploys a helper appliance to the target host to facilitate the mount, and in v13 a dedicated Linux mount server can serve this role across the infrastructure.

> **Note**
>
> When using enhanced role-based access control (RBAC) in v13, restore operators can be scoped to specific backups and allowed restore destinations, ensuring users only recover data within their authorized scope.

Volume-level restore

Volume-level restore is used when an entire disk volume has become corrupted, lost, or needs to be rolled back to a previous state, but a full machine restore is unnecessary or undesirable. This option is available for both virtual machine backups and physical machine backups created by Veeam agents.

The volume level restore wizard guides the administrator through the following steps:

1. Launch the volume level restore wizard from **Home** | **Backups** (right-click | **Restore virtual disks** | **VMware vSphere...**)

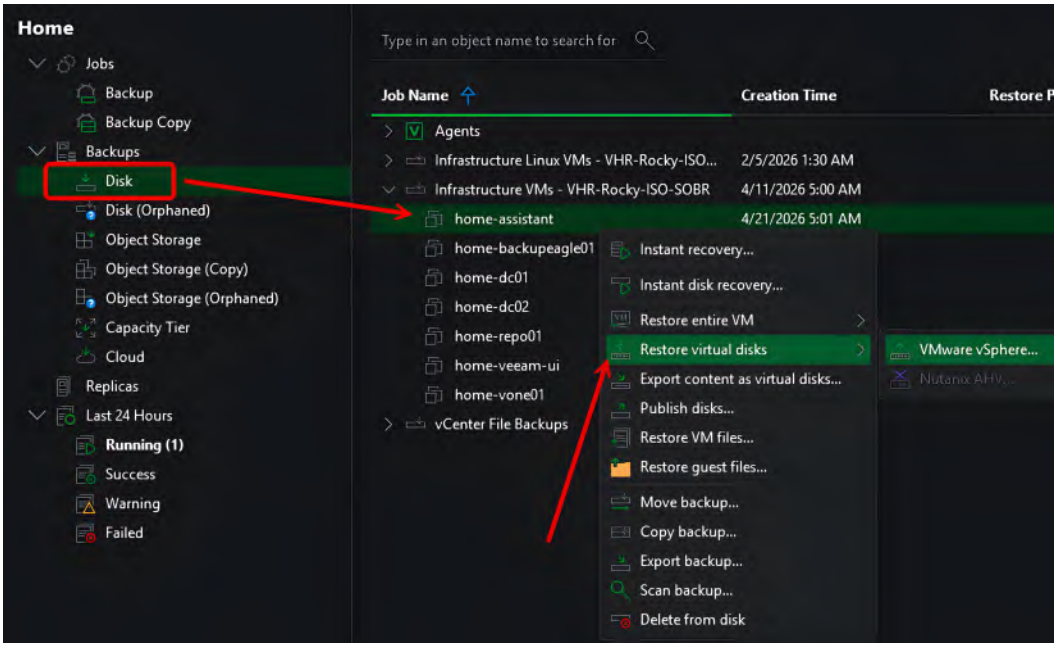


Figure 12.7: Restore virtual disk option

Technet24

2. Choose the required restore point.

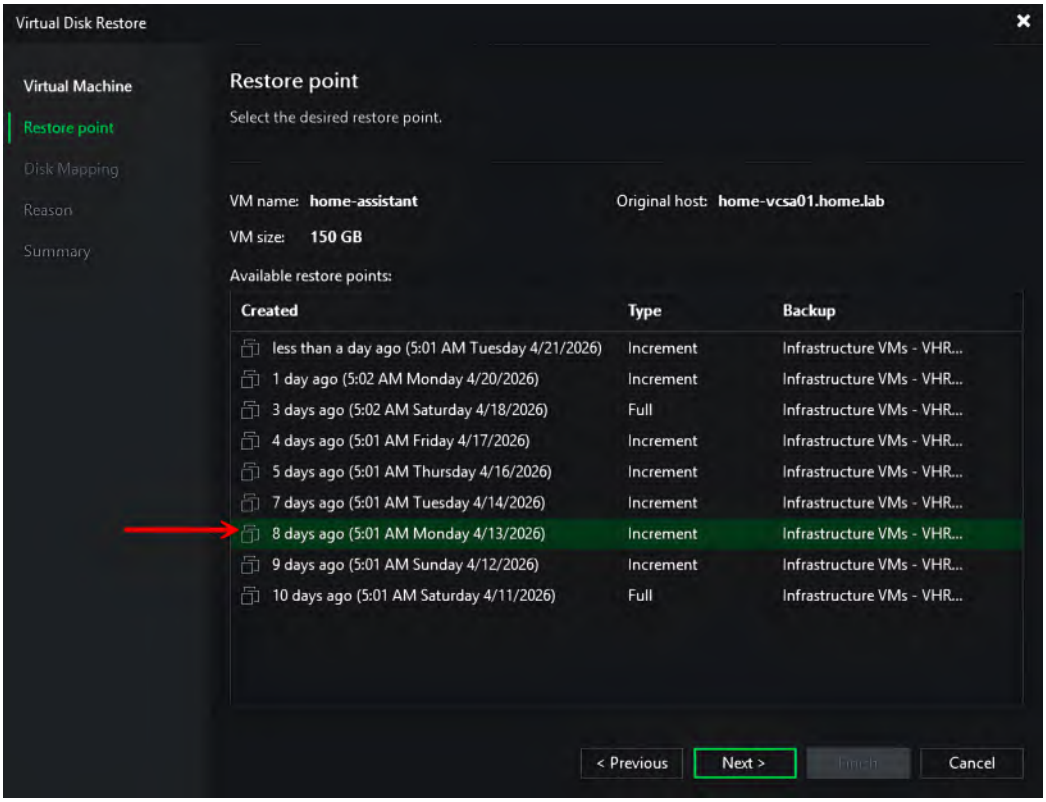


Figure 12.8: Restore point selection

3. Map the restored disk(s). By default, Veeam maps to the original location. Custom disk mapping is available for alternate targets.

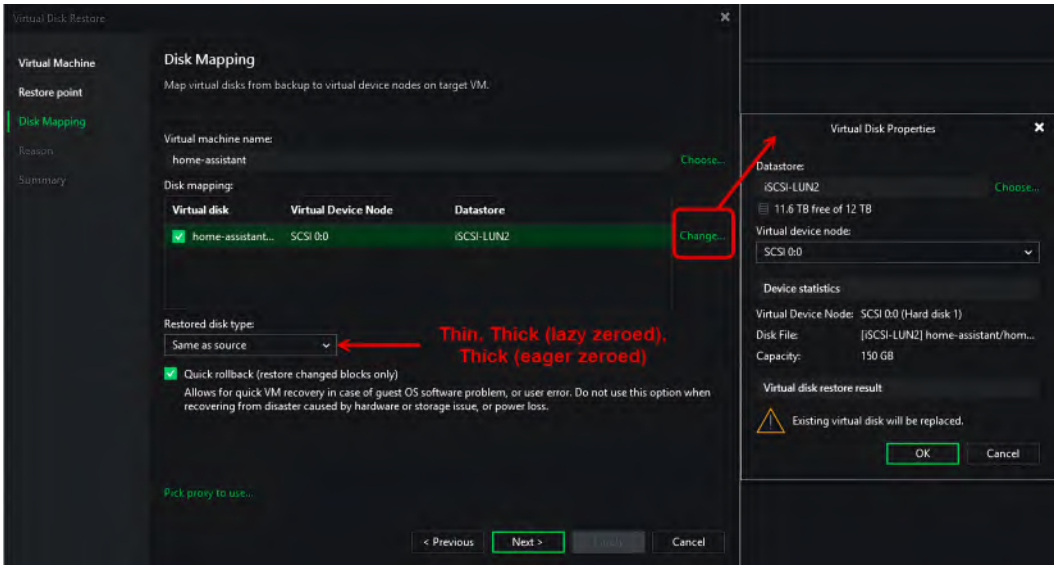


Figure 12.9: Disk mapping options

4. Optionally resize volumes if the target disk is larger or smaller than the original.
5. Configure secure restore settings (optional: see *step 4* of the restore wizard)
6. Specify a restore reason for audit logging.
7. Review the summary and launch the restore.

Volume-level restore supports flexible disk mapping options. When restoring to an alternate location, administrators can choose **Apply Backup Layout** (use the original disk configuration), **Apply Disk Layout** (copy the partition scheme from another disk), or **Erase** (start fresh). The target disk must be online and initialized prior to the restore.

Note

Important:

When restoring a system or boot volume, the machine must be booted from Veeam recovery media to perform the restore. Restoring the running system's active volumes is not supported and will require the machine to be taken offline or a Bare Metal Recovery to be performed instead.

Virtual disk restore (instant disk recovery)

Virtual disk restore (also referred to as **instant disk recovery**) allows administrators to restore individual virtual hard disks (VMDKs for vSphere, VHDs/VHDXs for Hyper-V) from a

Technet24

backup without needing to restore the entire virtual machine. This is particularly valuable when a VM has multiple disks and only a single disk has been lost or corrupted.

Unlike FLR which works at the file level, virtual disk restore operates at the disk level. The restored disk can be attached to the original VM, a different VM, or exported as a standalone virtual disk file. This makes it a powerful option for scenarios such as recovering a secondary data disk attached to a database server, restoring a corrupted application volume, or providing a test copy of a disk to an alternative VM without impacting production.

Restoring a virtual disk

Follow these steps to restore a virtual disk:

1. In the Veeam Backup & Replication console, go to **Home** | **Backups**.
2. Right-click the target VM and select **Instant disk recovery...**

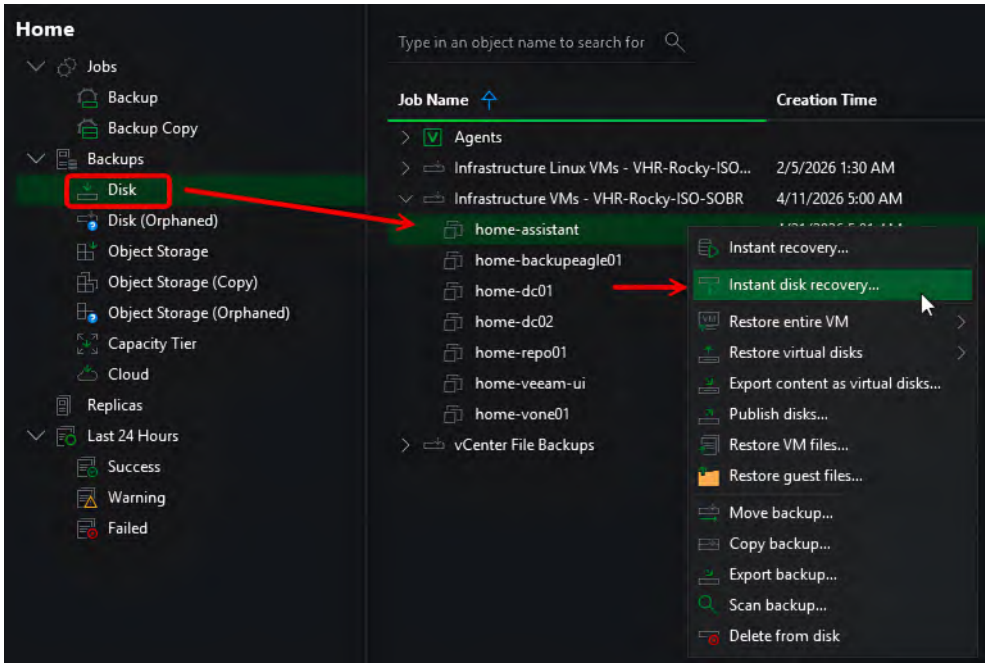


Figure 12.10: Instant disk recovery option

3. Select the appropriate restore point.

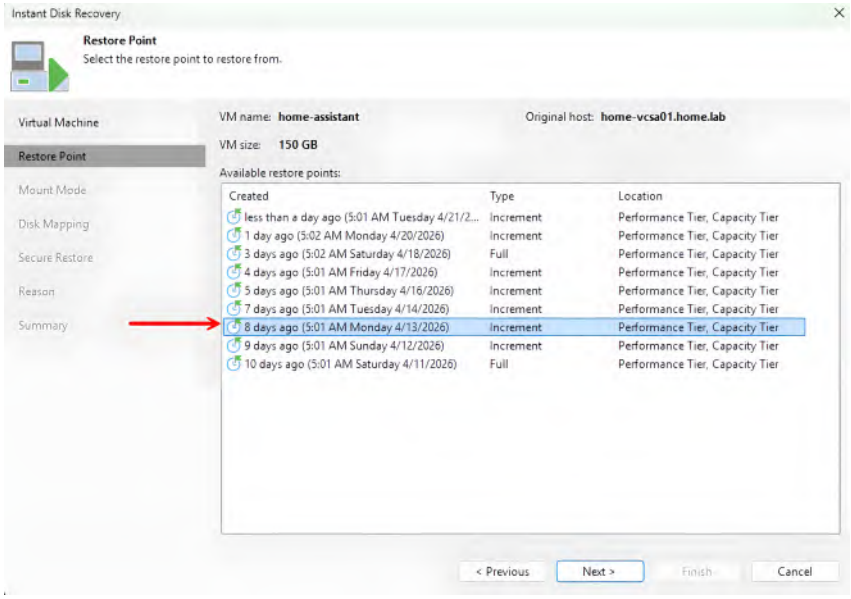


Figure 12.11: Restore point selection

4. Choose the mount mode for the virtual disk.

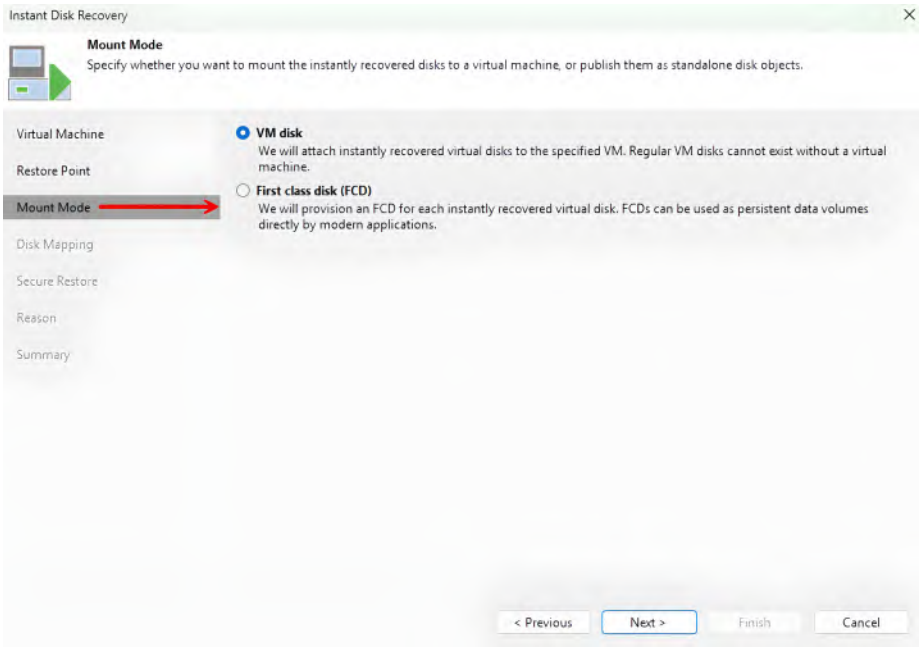


Figure 12.12: Mount mode for the virtual disk recovery

Technet24

5. Specify the disk mapping for the virtual disk.

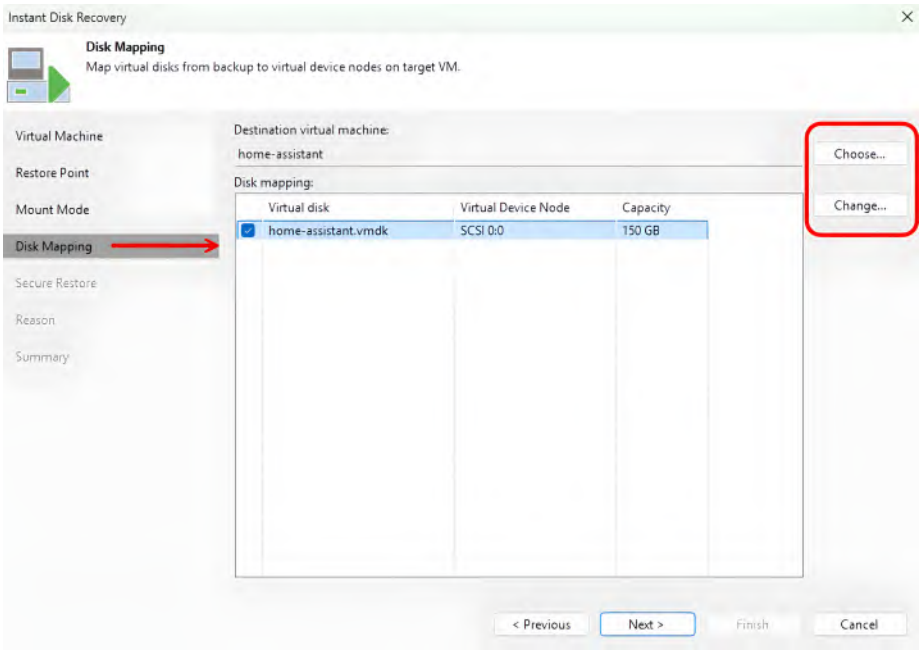


Figure 12.13: Disk mapping options for disk recovery

6. Select the **Secure Restore** options.

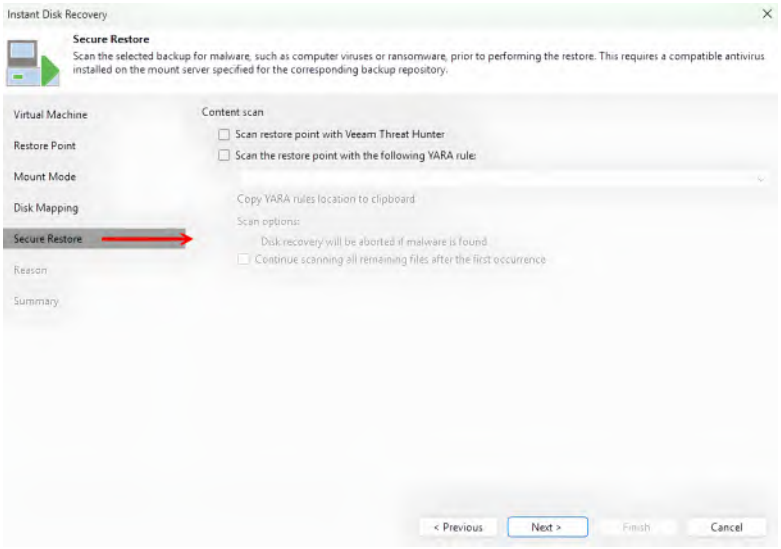


Figure 12.14: Secure restore options

7. Specify a restore reason for audit logging.
8. Review the summary and launch the restore. Veeam will write the restored disk directly to the specified destination.

Tip

Virtual disk restore is ideal for recovering a secondary data disk while the VM remains online. The production VM does not need to be powered off if the disk being restored is not the active system disk.

Application item restore — Veeam explorers

Veeam explorers provide granular, application-aware item-level recovery directly from image-level backups. When a backup job has application-aware processing enabled, Veeam captures a consistent, transactionally-aware snapshot of the application data. Veeam explorers then mount this backup content and present the application's data objects through a dedicated browser interface, enabling point-in-time recovery of individual items without restoring the entire virtual machine.

In Veeam Backup & Replication v13, the following Veeam Explorers are included:

- **Veeam explorer for Microsoft Active Directory:** Restore individual AD objects (users, groups, OUs, GPOs, contacts) or entire containers to the original domain or an alternate domain. Supports comparison of changed attributes since the backup was taken, enabling highly targeted restores.
- **Veeam explorer for Microsoft SQL server:** Perform point-in-time database restores, restore individual databases to a production or staging SQL server, or export databases as `.bak` files. Transaction log shipping is supported when log backup is configured, enabling sub-job-interval recovery points.
- **Veeam explorer for Microsoft Exchange:** Restore individual mailboxes, emails, calendar items, contacts, and public folders. Exchange server subscription edition (SE) is fully supported in v13. Secure LDAPS protocol can now be used instead of LDAP for exchange explorer connectivity, improving security posture.
- **Veeam explorer for Microsoft SharePoint:** Recover SharePoint sites, document libraries, lists, and individual items. SharePoint subscription edition 25 H1 is now fully supported in v13, including application-aware processing and item-level recovery.

Technet24

- **Veeam explorer for Oracle:** Restore Oracle databases to any point in time, recover individual tablespaces or data files, and export databases. Supports both host-based and agent-based backups when Oracle RMAN integration is active.
- **Veeam explorer for PostgreSQL:** Introduced in recent Veeam versions and fully supported in v13, this Explorer provides granular database-level and point-in-time recovery for PostgreSQL instances. v13 also supports application-aware processing for PostgreSQL on Proxmox VE in version 13.0.1.

Launching a Veeam explorer

To open a Veeam explorer, right-click the backed-up machine in **Home | Backups**, select **Restore Application Items**, and choose the target application (e.g., Microsoft SQL server). The wizard prompts the restore point, then automatically mounts the backup and launches the corresponding explorer. A staging server may be required for SQL and Oracle restores to temporarily host the database during the item extraction process.

> **Note**
>
> Application-aware processing must be enabled in the backup job configuration for Veeam explorers to function. Without it, Veeam cannot guarantee transactional consistency, and the explorer may not be able to mount or parse the application data correctly.

Export backup

The **export backup** feature allows administrators to synthesize a complete, self-contained full backup file (`.VBK`) from any restore point in an existing backup chain, regardless of whether that point is stored as a full backup, incremental, or reverse-incremental backup file. Veeam reads all dependent backup files, merges the data, and writes an independent `.VBK` file to a specified destination.

Common use cases for export backup include the following:

- **Legal hold and archiving**: Preserving a specific restore point as a standalone file to prevent it from being removed by the job's retention policy.
- **Offsite portability**: Providing a single portable backup file on removable media (USB drive, tape) for offsite storage or handoff to another team.
- **Cross-environment recovery**: Exporting a restore point so it can be imported into a different Veeam installation or backup repository.
- **Compliance and audit**: Retaining point-in-time snapshots of critical systems as immutable evidence records.

To run an export backup, right-click a backup job or individual VM backup in **Home | Backups**, select **Export Backup**. Then choose the desired restore point, specify the destination path and any retention settings for the exported file and complete the wizard. The resulting .VBK file is fully portable and can be imported into any Veeam repository using the **Import Backup** function.

Tip

Export backup is also useful for creating a known-good recovery point before a major change such as an application upgrade or OS patching cycle, ensuring a clean restore point is available independent of the regular backup chain retention schedule.

Secure restore

Secure restore is a critical cyber-resilience feature that allows administrators to scan backup data for malware before it is restored to the production environment. This prevents the re-introduction of ransomware, viruses, or other malicious software that may have been active at the time the backup was taken—a scenario that is particularly dangerous during ransomware recovery where the initial infection point may be unknown.

In Veeam Backup & Replication v13, secure restore integrates tightly with the Veeam threat hunter engine and third-party antivirus solutions. The malware scan is performed before the restored data is committed to the target, ensuring that potentially infected data never reaches production systems. v13 extends Veeam threat hunter scanning to Linux machines and cloud workloads (AWS, Azure, Google Cloud), significantly broadening the protection surface.

Securing restore integration with recovery

Secure restore can be enabled within the following recovery wizards:

- Entire VM restore
- Volume-level restore
- Virtual disk restore
- Instant recovery (as a pre-recovery scan option)

When secure restore is enabled, the wizard presents a step to configure the antivirus scanner. The administrator can choose between Veeam threat hunter (signature-based, built-in) or a custom AV solution deployed on the mount/staging server. If malware is detected during the scan, the administrator is presented with options to abort the restore, continue with the infected data quarantined, or proceed anyway. Each action is logged for audit purposes.

Technet24

Note

Best practice:

Always enable secure restore when recovering data following a suspected or confirmed malware or ransomware incident. Scanning multiple restore points to identify a clean recovery point prior to the infection is strongly recommended before committing to production.

Instant cloud recovery to Microsoft Azure (v13)

Instant recovery to Microsoft Azure is a major new capability which extends Veeam's proven instant recovery technology, which was previously limited to on-premises hypervisors, to the public cloud, enabling any Windows or Linux machine protected by Veeam to be recovered as a native Azure virtual machine in under five minutes, including Azure VM provisioning and boot to the OS logon screen.

This capability is particularly compelling for disaster recovery scenarios where on-premises infrastructure is unavailable or compromised. Unlike traditional cloud failover approaches that require significant pre-configuration or replication pipelines, instant recovery to Azure works directly from existing Veeam backups, including backups stored in Veeam cloud vault (Veeam vault). The recovered Azure VMs can, in some I/O scenarios, perform comparably to or better than natively provisioned Azure VMs due to Veeam's optimized I/O stack.

Use cases for instant cloud recovery include the following:

- **Disaster recovery (DR) to Azure:** Fail over workloads to Azure when the primary site is down, with RTOs under five minutes.
- **Clean room recovery:** Spin up an isolated Azure environment for forensic investigation or malware scanning without impacting production.
- **Data Classification and Analysis:** Recover workloads to Azure temporarily for data discovery, compliance scanning, or analytics workloads.
- **Testing and validation:** Verify restore integrity and application functionality in Azure without consuming on-premises resources.

Note

Instant recovery to Microsoft Azure requires appropriate Azure subscription credentials to be configured in Veeam and sufficient Azure quota for the target VM size. Costs associated with running recovered Azure VMs are billed to the Azure subscription and

should be monitored. For long-term DR, consider using Azure as a replica target rather than running workloads continuously from backup.

Bare metal recovery (BMR)

Bare metal recovery (**BMR**) is used to restore a physical or virtual machine to a completely different hardware configuration, or to restore a system whose operating system is no longer bootable. BMR is typically performed by booting the target machine from Veeam recovery media (an ISO-based bootable image) and connecting to a Veeam backup repository to restore the entire machine image.

In Veeam v13, ACL (access control list) preservation has been improved for Linux systems during BMR when using the NFSv4 protocol. When NFSv4 is unavailable, the process automatically falls back to NFSv3, though ACL data will not be included in that scenario. Additionally, v13 adds support for connecting to a backup server in read-only mode as an additional DR connection option, providing flexibility in constrained recovery situations.

BMR process overview

The following steps outline the process for conducting a bare metal recovery using the Veeam recovery media:

1. Create Veeam recovery media (ISO) from the Veeam Backup & Replication console by navigating to **Home** | **Backup Infrastructure** | **Create Recovery Media**.
2. Boot the target machine from the recovery media (USB, CD, or virtual ISO mount via IPMI/iDRAC/iLO)
3. The Veeam recovery wizard launches automatically. Select the network settings to connect to the backup repository.
4. Authenticate with the Veeam backup repository and select the machine backup to restore.
5. Choose the desired restore point from the available chain.
6. Select the restore mode — **Entire Computer** for a full BMR or choose specific volumes to restore.
7. Map the source disks from the backup to the target physical disks. Adjust disk layout as needed.
8. Review the summary and initiate the restore. The process duration depends on the backup size and network throughput.
9. Once complete, the machine reboots from the restored system.

Technet24

> **Note**
>
> **Important:**
>
> Recovery media should be recreated and tested after each Veeam Backup & Replication upgrade to ensure compatibility with the current software version. It is also recommended to test BMR procedures in a lab environment at least quarterly as part of a formal DR testing program.

Quick reference for recovery options

The following table summarizes each recovery option, its primary use case, and the recovery scope to assist in selecting the right restore path for a given scenario:

Recovery option	Primary use case	Scope	VM online during restore?
File-level recovery	Recover deleted/ corrupted files or folders	File/Folder	Yes
Volume-level restore	Restore a corrupted or lost volume	Disk volume	No (for system volumes)
Virtual disk restore	Recover a specific VM virtual disk	Virtual disk (VMDK/ VHDX)	Yes (non-system disks)
Application item restore (Veeam explorers)	Recover individual app objects (AD users, SQL DBs, emails, etc.)	Application item	Yes
Export backup	Archive, legal hold, or portability of a restore point	Full backup chain	N/A (offline operation)
Secure restore	Malware-safe recovery after a cyber event	VM, volume, or disk	Depends on restore type

Recovery option	Primary use case	Scope	VM online during restore?
Instant cloud recovery (Azure)	Rapid DR or clean room recovery in Microsoft Azure	Full machine	N/A (cloud-based)
Bare metal recovery	Full physical/virtual machine restore to new hardware	Full machine	No

Table 12.7: Available recovery options

Choosing the right recovery option

Selecting the appropriate recovery method is as important as having the backup itself. For day-to-day incidents such as accidental file deletion or individual email recovery, file-level recovery, and Veeam explorers provide fast, targeted restoration with minimal impact on production systems. For more severe scenarios involving corrupted volumes or failed system disks, volume-level restore, and bare metal recovery provide the necessary depth of recovery.

In ransomware or security incident scenarios, secure restore should be incorporated into every recovery workflow to prevent re-infection. Veeam v13's integration of Veeam threat hunter across physical, virtual, Linux, and cloud workloads provides a unified scanning capability that eliminates gaps in the security recovery process. For organizations with cloud DR requirements, instant recovery to Microsoft Azure in v13 offers a compelling option to minimize RTO without requiring ongoing VM replication costs.

Administrators should document their recovery procedures for each option, test them regularly as part of a formal DR testing program, and leverage Veeam's RBAC to ensure the right personnel have access to the right recovery tools, without exposing the backup infrastructure to unnecessary risk.

Summary

This chapter has provided you with many of the recovery options available in Veeam Backup & Replication v13. We discussed instant recovery to Microsoft Azure, what it is, and how it is used. Then, we investigated Universal CDP and how it can assist with recovery of services. Lastly, we examined the other recovery options available within the platform and how best to use them.

Technet24

This chapter concludes the fourth edition of the *Mastering Veeam Backup & Replication v13* book, and I want to thank you for reading it. I hope this book's content will benefit you and help you configure your environment based on the topics covered, including many of the newest enhancements and features added to the Veeam products.

Further reading

Instant Recovery to Azure - `https://helpcenter.veeam.com/docs/vbr/userguide/instant_recovery_to_azure.html?ver=13`

13

Unlock Your Exclusive Benefits

Your copy of this book includes the following exclusive benefits:







Technet24

Follow the guide below to unlock them. The process takes only a few minutes and needs to be completed once.

Unlock this Book's Free Benefits in 3 Easy Steps

Step 1

Keep your purchase invoice ready for *Step 3*. If you have a physical copy, scan it using your phone and save it as a PDF, JPG, or PNG.

For more help on finding your invoice, visit `https://www.packtpub.com/en-us/unlock?step=1`.

> **Note**
>
> If you bought this book directly from Packt, no invoice is required. After *Step 2*, you can access your exclusive content right away.

Step 2

Scan the QR code or go to `packtpub.com/unlock`.

On the page that opens (similar to *Figure 13.1* on desktop), search for this book by name and select the correct edition.

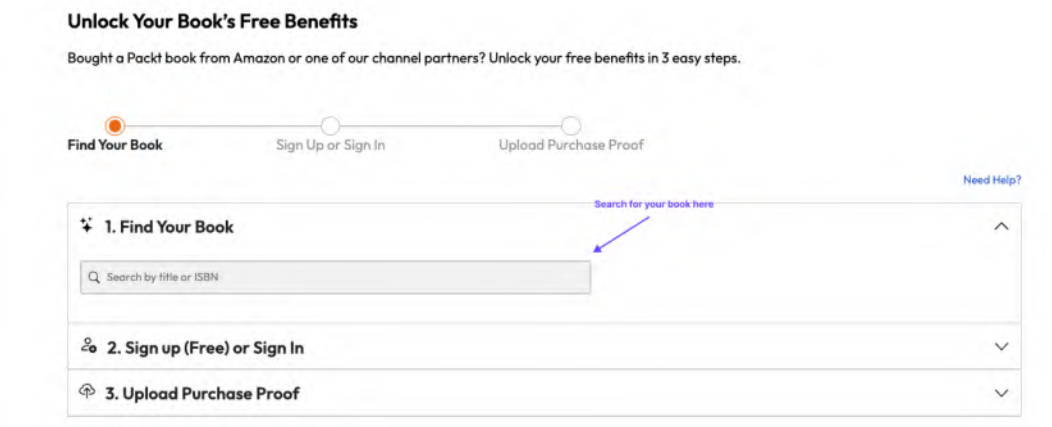


Figure 13.1: Packt unlock landing page on desktop

Technet24

Step 3

After selecting your book, sign in to your Packt account or create one for free. Then upload your invoice (PDF, PNG, or JPG, up to 10 MB). Follow the on-screen instructions to finish the process.

Need Help

If you get stuck and need help, visit `https://www.packtpub.com/unlock-benefits/help` for a detailed FAQ on how to find your invoices and more. This QR code will take you to the help page.



Note

If you are still facing issues, reach out to `customercare@packt.com`.



packt.com

Subscribe to our online digital library for full access to over 7,000 books and videos, as well as industry leading tools to help you plan your personal development and advance your career. For more information, please visit our website.

Why subscribe?

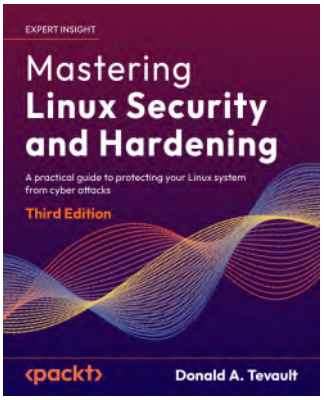
- Spend less time learning and more time coding with practical eBooks and Videos from over 4,000 industry professionals
- Improve your learning with Skill Plans built especially for you
- Get a free eBook or video every month
- Fully searchable for easy access to vital information
- Copy and paste, print, and bookmark content

At www.packt.com, you can also read a collection of free technical articles, sign up for a range of free newsletters, and receive exclusive discounts and offers on Packt books and eBooks.

Technet24

Other Books You May Enjoy

If you enjoyed this book, you may be interested in these other books by Packt:

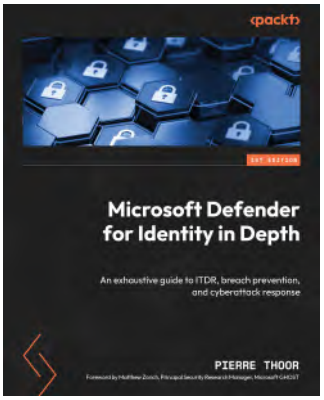


Mastering Linux Security and Hardening

Donald A. Tevault

ISBN: 9781837630516

- Prevent malicious actors from compromising a production Linux system
- Leverage additional features and capabilities of Linux in this new version
- Use locked-down home directories and strong passwords to create user accounts
- Prevent unauthorized people from breaking into a Linux system
- Configure file and directory permissions to protect sensitive data
- Harden the Secure Shell service in order to prevent break-ins and data loss
- Apply security templates and set up auditing



Microsoft Defender for Identity in Depth

Pierre Thoor

ISBN: 9781835884485

- Ensure a secure and efficient MDI setup for peak defense capabilities
- Unlock automation with PowerShell scripting magic
- Seamlessly blend MDI with AD CS, AD FS, and Entra Connect
- Expand MDI's reach and impact through APIs
- Pioneer advanced threat hunting with KQL expertise in Defender XDR
- Craft expert, strategic responses to security alerts
- Optimize action accounts for maximum agility
- Establish a robust and rigid ITDR framework

Packt is searching for authors like you

If you're interested in becoming an author for Packt, please visit `authors.packtpub.com` and apply today. We have worked with thousands of developers and tech professionals, just like you, to help them share their insight with the global tech community. You can make a general application, apply for a specific hot topic that we are recruiting an author for, or submit your own idea.

Technet24

Share your thoughts

Now you've finished *Mastering Veeam Backup Replication*, we'd love to hear your thoughts! Scan the QR code below to go straight to the Amazon review page for this book and share your feedback or leave a review on the site that you purchased it from.



https://packt.link/r/1806388898

Your review is important to us and the tech community and will help us make sure we're delivering excellent quality content.

Index

Symbols

64-bit architecture **289, 290**
- cross-platform architecture, for future 296
- data mover improvements 295
- instant recovery engine 295
- minimum supported operating systems 290

A

AI-powered solution **386**
- architecture and integration 386
- deep data analysis agent 387
- malware threat analysis agent 387
- operating modes 387

Active Directory **177**

Active Directory integration, Veeam Backup & Replication **177**
- authentication model 178
- best practices 180, 181
- integration issues, troubleshooting 181, 182
- prerequisites 178
- role-based access control, configuring with AD groups 179
- security considerations 180, 181
- service account configuration and delegation 180
- with Veeam software appliance 182, 183

Active Directory integration, Veeam Backup & Replication v13
- agent-based protection groups 193
- authentication 191
- best practices 183
- deployment model comparison 188, 189
- deployment options 185
- domain 184
- domain placement strategy 184
- enhanced role-based access control 191, 192
- forest 184
- forest trust configuration 189, 190
- identity management 191
- management domain deployment 187
- security recommendations 193, 194

Active Directory integration, Veeam Backup and Replication v13
- Windows workgroup deployment 186

Active Directory security, Veeam Backup & Replication v13
- deployment models 195
- DNS considerations 203
- management forest 195, 196
- multi-factor authentication (MFA) 201
- network segmentation 203
- production domain 198
- service accounts, securing 199
- workgroup deployment 197, 198

Azure Blob, immutability improvements **341**
- container-level policy validation 341
- immutability policy renewal 341

Azure account
- requirements 400

Azure subscription
- requirements 400

active node **231**

air-gapped environments
- updates, configuring for 81

alarm engine integration capabilities **348**
- application backup policy state alarm (v13.0.1) 349
- Coveware recon integration alarms (v13.0.1) 349
- VBR authentication anomaly alarms (v13.0.1 Patch 1) 349
- VBR high availability health alarms (v13.0.1) 349

application backup policy state alarm (v13.0.1) **349**

application drive **5**

application item restore **415, 425**

archive tier improvements **338**
- archive retrieval cost estimation 339

Technet24

selective archive tiering, by retention policy 339
authentication modernization 293
automated patching 62
automated updates 62
automatic failover 233
characteristics 233

B

BLAKE3 hashing 294
Backup from Storage Snapshots 21
backup administrator 173
best practices 175
capacity planning and forecasting 174
compliance and reporting 174
daily operations, simplifying 173, 174
knowledge assistance 175
practical considerations 175, 176
backup data immutability 337
backup repository operational changes
scaling out 342
backup server 208
bare metal recovery (BMR) 415, 429
process overview 429
block generation 112
benefits 113
periods by storage type 113
storage capacity, planning 114
working 112

C

CDP coordinator service 245
CDP health
monitoring 413
troubleshooting 413, 414
CDP network resilience
best practices 276, 277
CDP policy 252
creating 262
guest processing, configuring 268
network mapping, configuring 268
proxies, assigning 265
RPO settings, defining 266, 267
source VMs, selecting 263
summary and application 269, 270
target, configuring 264
wizard, launching 262, 263
CDP proxy 407
Configuration Backups 18
Console Local Cache 210
Coveware recon integration 391
Coveware recon integration alarms (v13.0.1) 349
Coveware recon scanner 375
Credentials Manager 355
catalog drive 5
centralized update management
benefits 82, 83
centralized updates, with Veeam updater 306
architecture 306
automatic updates, configuring 307
components 306
updates, monitoring 310
updates, troubleshooting 310
cloud workload monitoring expansion (v13.0.1) 353, 354
cluster DNS Name 210
compliance mode 340
container-level policy validation 341
continuous data protection (CDP) 405
copy backups 131
custom role
creating 71
custom role-based access control (RBAC) 312
architecture 312
assignments 313
built-in roles, versus custom roles 313
custom roles, creating 314 – 318
permission categories 319, 320
permissions 312
role assignment process 320, 321
roles 312
scope inheritance and precedence 321, 322

D

DISA STIG compliance 59, 107, 108
security controls 108
Dashboard Catalog
redesigning 348
Data Platform Scorecard 347

Default Backup Repository **18**
Defense Information Systems Agency (DISA) **59**
Direct NFS **21**
Direct Storage Access **21**
data collection settings **355**
data resilience daily summary **156**
data resilience summary (v13.0.1) **357 – 360, 388**
database maintenance
configuring 329
connection pooling 330
index bloat 330
session history table growth 329
deep data analysis agent **156, 387**
deep-data analysis agent (v13.0.1) **357**
direct backup, to object storage **339**
immutable backup chain integrity verification 340
synthetic full backup support 339
disaster recovery as a service (DRaaS) **285**
Veeam cloud connect 285
dispatch layer improvements **333**
proxy selection algorithm optimization 333
repository queue management 333
resource utilization tracking, decoupled from scheduling 333
domain **184**
dual-database architecture **363**
administrators, implications 365
database (PostgreSQL), reporting 364
database (SQL server), monitoring 364

E

enhanced Role-Based access control (RBAC) **70**
security benefits 74
enterprise application plug-in monitoring (v13.0.1) **353**
etcd **232**
export backup feature **415, 426**
use cases 426

F

Fast Clone
working 114, 115
FastSCP **120, 136**
File Backup Proxy **17**
failover **226**
best practices, for reliability and effectiveness 240
issues and troubleshooting 238, 239
versus switchover 227
failover and recovery enhancements
file-level restore, from CDP replicas 284
I/O throughput visualization, during failover 283
planned CDP failover 284
failover events
monitoring 235
file-level recovery (FLR) **414, 415**
Windows, versus Linux 418
working 416 – 418
forest **184**

G

Google remote procedure call (gRPC) **291**
impact, on network architecture and firewall rules 292
Grandfather-Father-Son (GFS) backups
and immutability 115
global update
settings 77 – 79
governance mode **340**
group managed service accounts (gMSA) **199**
prerequisites 200

H

Heatmaps **385**
High Availability (HA) **63**
benefits 67
for disaster recovery 208
licensing and version requirements 66, 67
High Availability (HA) cluster **208, 213**
architecture 209
assembling, via VBR console 221 – 223
assembly 215
automatic updates 228, 229
benefits 229, 230
Cluster DNS name 210, 214

Technet24

Cluster IP address (Virtual IP) 64, 214
configuring 218
considerations 230
Console Local Cache 210
core components 209, 210
data replication 210
data synchronization 210, 215
disassembling 228
disaster recovery design patterns 211, 212
DNS configuration 217
emergency failover 65, 66
enabling, on both nodes 218 – 220
failover event sequence 232, 233
internal components 218
licensing 216
limitations 230
network configuration 217
planned switchover 65
PostgreSQL instance 209, 218
primary node 64, 214
rescanning 228
secondary node 64, 214
server requirements 216
split-brain prevention 211
status, monitoring 228
VBR console cache 218
veeamhasvc.service 209, 218
virtual cluster IP 209
WAL file replication 209
WAL file storage and resilience 216 – 218

Host Management Console (Port 10443) **60**

Hot-Add Mode **21**

hardened Linux repository immutability **342**
immutability period per backup job 342
immutability verification reports 342

hardened repository **302**
adding, to veeam backup and replication 104, 105
backup repository, configuring, with immutability 105
deploying 102
immutability 111, 112
immutability, considerations 106
immutability, working 112
installing 103, 104
system requirements 102

home section, Veeam console
Backups view 130, 131
Jobs view 129, 130
Last 24 Hours view 132, 133
main menu 124 – 129

I

I/O journal **252**
architecture and storage mechanics 278
scalability benchmarks 279
scalability enhancements 278

I/O journal engine **407**
improvements 407

ISO deployment method **49**

Incident API Operator **71**

immutability configurations
best practices 116

immutability enhancements **340**
lock period alignment, with retention 340
S3 object lock improvements 340
status, in VBR console 341

immutability period per backup job **342**

immutability policy renewal **341**

immutability verification reports **342**

imported backups **131**

infrastructure monitoring updates **352**

instant recovery to Microsoft Azure **395, 396, 401 – 403, 415, 428**
helper appliance templates 397, 398
innovations 397
mass recovery considerations 404, 405
migrate to production (switchover) 403
network considerations 401
recovery process flow 398
stop publishing (discard) 404
supported backup sources 398
supported repository types 399

interactive dashboards and widgets **384**
custom dashboard creation 386
Heatmaps 385
highlights, predefining 385

J

JeOS ISO **89**
backup infrastructure lockdown mode 109
built-in immutability 109

designated security officer role 109
minimized attack surface 110
operational implications 111
privilege separation 108
ransomware protection 109
scheduling and control updates 110
security benefits 110
SSH disabled, by default 110
veeam updater, for automated security updates 109
Zero Trust architecture 108

Jobs Calendar view 354, 355

Just enough operating system (JeOS) 88
appliances, managing 96, 97
deployment, benefits 89, 90
deployment, system requirements 91
versus Veeam software appliance 88

Just enough operating system (JeOS) architecture 59
characteristics 59

job manager 331
architecture 331, 332
scalability health, monitoring 336

job manager scalability
considerations 336

jump box 122

K

Kerberos
versus NTLM authentication 192, 193

keepalived 232

L

Linux Backup Repositories 28

Linux infrastructure roles, in v13 298
additional roles 299
gateway servers 298
guest interaction proxy (GIP) 299
mount server 298

Linux-only backup infrastructure 296, 297
certificate-based authentication 301
core architecture components 302
designing 302
infrastructure roles, in v13 298
migration considerations 304
reduced port requirements 301
role placement guidelines 302
security benefits 305
Veeam Infrastructure Appliance (VIA) 299, 300
Windows console considerations 303

large-scale deployment guidance 335
backup window design 335
VBR server sizing 335

least-privilege access policies 62

limitations page, Veeam Backup and Replication website
reference link 30

local journal spooling mechanism 408

long-term restore points 246

M

Malware Detections widget 347

Microsoft Active Directory (AD) 177

Microsoft Hyper-V 396

malware deep analysis agent 172

malware threat analysis agent 155

malware threat analysis agent (v13.0.1) 357

management domain
advantages 187, 188
disadvantages 188

management forest 195, 196

mandatory updates 76

manual failover 234

morning coffee report 156

multi-factor authentication (MFA) 61, 201
enabling, in Veeam console 201, 202
for Veeam software appliance host management UI 202

N

NTLM authentication
versus Kerberos 192, 193

No Per-VM Selection 26

Nutanix AHV 396

network mode 21

network resilience, improving for CDP 271
failure modes 272
network throttling rules, configuring 273

Technet24

network traffic encryption and compression 273
preferred networks 274
WAN accelerator integration 274
network resilience, with CDP
CDP health, monitoring 276
continuous data protection 250
improving 271
proxy failover 275, 276
redundant network paths 275, 276
RPO compliance 276
source proxy cache configuration 275

O

OS drive 4
OVA deployment method 49
Object Storage 28
object storage and immutability
considerations 343
object storage enhancements 337
capacity tier offload performance 337, 338
multipart upload reliability 338
object storage integration 337
one-way forest trust 190
operating modes, Veeam Intelligence 170
advanced 171
agent-based retrieval 171
basic 171
code generation and analysis 172
data access, in advanced mode 172
disabled 170, 171
malware deep analysis agent 172
multi-agent architecture 171
optional updates 76
orphaned backups 131

P

Patroni 231
Patroni logs 235
Per-VM Backup Files 25
PostgreSQL 326
considerations 330
limitations 330
upgrading, from previous versions 329
PostgreSQL 17
key improvements, relevant to VBR 326
PostgreSQL 17.6 295, 326
supported deployment configurations 327
VBR v13.0.1, installation 327, 328
PostgreSQL instance 209
PowerShell integration
for RBAC management 323, 324
Proxy Servers 15 – 17, 20
configuring 21 – 23
optimizing 21 – 23
passive nodes 231
permissions, Veeam Backup & Replication v13
reference link 10
planned failover 410
planned switchover
workflow 234
platform-independent I/O journal engine 406
point-in-time failover 409, 410
post-failover operations 236
active node, validating 236
failed node, recovering 236
post-failover recovery 227
production domain 198

R

RBAC management
PowerShell integration 323, 324
REST API authentication 374
RPO anomalies widget 347
ReFS/XFS 25
Repository Server 15, 18
read-only mode 211
read-write mode 211
real-world RBAC design patterns 322
auditing and compliance 323
geographic segmentation 322
separation of duties 322
tiered administration 322
recovery operations, Universal CDP
failback and reverse protection 410
planned failover 410
point-in-time failover 409, 410
recovery options 414
application item restore 415, 425

bare metal recovery (BMR) 415, 429
export backup 415
file-level recovery (FLR) 414 – 418
instant cloud recovery to Microsoft Azure 415, 428
reference 430, 431
secure restore 415, 427
selecting 431
virtual disk restore 415, 421
volume-level restore 415, 419 – 421

redesigned report catalog 351, 352

reporting engine 349, 388
database configuration, reporting 390
export format changes 350
HTML5-based interactive reports 389
key capabilities 350
PostgreSQL reporting database 388, 389
redesigned report catalog 351, 352, 389, 390

repository servers
setting up 25, 26

reverse CDP protection 410

role configurations
examples 72, 73

S

SAML 2.0 single sign-on 191
benefits 191

SLA Compliance Overview 348

SOBR management
reference link 38

SOBR rebalance 38

Scale-Out Backup Repository (SOBR) 27 – 30
archive tier 30
capacity tier 30 – 33
creating 33 – 36
Enterprise 29
Enterprise Plus & VUL (Veeam Universal License) 29
performance tier 30, 31
reference link 37

Security Administrator 71

Security Officer 71
authorization 58
capabilities 57, 58

Security Technical Implementation Guide (STIG) 59

ServiceNow integration enhancements 391

Shared Folder 28

SureBackup 135

scheduling engine improvements 332
after-job chaining 332, 333
parallel schedule evaluation 332
schedule condition caching 332

scheduling tick 331

secure remote access 59, 60

secure restore 415, 427
enabling 427

security features
summary 63

security hardening 59

security officer role 61

service accounts
securing 199

session state management 334
reduced memory footprint per active session 334
session history write optimization 334

short-term restore points 246

source CDP proxy 251

speech-to-text 158

standalone hosts 282

standard service account hardening 200

switchover 65, 224, 225, 234
versus failover 227

system logs 235

T

target CDP proxy 252

task limits
reference link 24 – 26

tasks 22, 23

testing cadence 238

threat center dashboard 346
Data Platform Scorecard 347
Malware Detections widget 347
RPO anomalies widget 347
SLA Compliance Overview 348

threat center integration and security reporting 391
Coveware recon integration 391

Technet24

ServiceNow integration enhancements 391
VSA monitoring 391

transport mode 21

U

UEFI Firmware Settings 50

Universal CDP 244, 281, 406
advantages 406
agent-based architecture 281, 282
CDP proxy 407
CDP target 408
constraints 248
coordinator service 245
data flow 283
deployment overview 247, 248
features 248
I/O journal continuity 408, 409
I/O journal engine 407
in context of broader recovery strategy 411
integration with Veeam cloud connect, for DRaaS 247
licensing 248
network resilience 408, 409
recovery operations 409
replica datastores 408
replication targets 246, 247, 282
RPO, balancing against infrastructure overhead 411
Veeam CDP agent service 244
Veeam CDP volume filter driver 244
versus traditional backup 412
volume filter driver 245
working 245

update architecture 75

update distribution model 306

update management
best practices 84

update types
global update 77 – 79
mandatory updates 76
optional updates 76
summary 82

updates
component-specific configuration 80
configuring, for air-gapped environments 81

V

VBR Web UI (Port 443) 60

VBR authentication anomaly alarms (v13.0.1 Patch 1) 349

VBR console integration 392

VBR high availability health alarms (v13.0.1) 349

VDP Premium Edition 67

VMware Backup Proxy 17

VMware I/O filter 251

VMware backup proxy deployment 98
infrastructure appliance, installing for proxy role 99, 100
system requirements 98, 99
VIA server, adding to Veeam backup and replication 100, 101
VMware backup proxy role, assigning 101

VMware backup proxy role
transport mode 101

VMware cloud director 246, 247, 282

VMware vCenter Credentials 15

VMware vSphere 246, 396

VMware vSphere clusters 282

VSA monitoring 391

VeeaMover 37

Veeam Backup & Replication 49, 345
installing 50 – 57
upgrading, to v13 38 – 43
Veeam intelligent diagnostics (VID) 356

Veeam Backup & Replication monitoring improvements 354
Credentials Manager 355
data collection settings 355
Jobs Calendar view 354, 355

Veeam Backup & Replication server 17

Veeam Backup & Replication user interface (UI)
history 119 – 121
updating 137, 138

Veeam Backup & Replication v13 289
installing 4 – 14
settings, configuring 15 – 19

Veeam Backup Administrator 70

Veeam Backup Enterprise Manager 50

Veeam Backup Operator 70

Veeam Backup Viewer **70**
Veeam Backup for Microsoft 365 **345**
Veeam CDP agent service **244, 281**
Veeam CDP volume filter driver **244, 245, 281**
long-term restore points 246
short-term restore points 246
Veeam Data Mover **21**
Veeam Data Platform **289, 345**
Veeam High Availability (HA) **207**
Veeam Infrastructure Appliance (VIA) **88, 299 – 302**
automated installation 93
configuration, completing 95
deploying 91
initial configuration wizard 94, 95
installation options 97, 98
role, selecting for deployment 92
system, booting 92
target system, preparing 91
VIA ISO, downloading 91
Veeam Infrastructure Appliance ISO
backup infrastructure roles 89
Veeam Installer Service **21**
Veeam Intelligence **123, 153 – 155, 345**
accelerated troubleshooting 159
accessing 161
best practices 168, 169
data resilience daily summary 156
deep data analysis agent 156
enhanced operational efficiency 161
enhancements, summary 158
improved cross-team communication 160
malware threat analysis agent 155
operating modes 170
practical use cases 161
prerequisites 162
proactive risk management 159, 160
reporting 158
root cause analysis 159
security considerations 169
simplified compliance and reporting 160
threat dwell time, reducing 160
transparent AI processing 158
troubleshooting 169
visualization 158
voice interaction capabilities, through speech-to-text 158
Veeam Intelligence, Windows console
accessing 166
chat interface 167
chat window 166
launching 166
settings, configuring 168
Veeam Intelligence, web UI
accessing 162
chat interface 163 – 165
configuring 165
launching 162
Veeam JeOS **88**
Veeam ONE
alarms 235
Veeam ONE v13 **66, 345, 346**
alarm engine integration capabilities 348
cloud workload monitoring expansion (v13.0.1) 353, 354
Dashboard Catalog, redesigning 348
enterprise application plug-in monitoring (v13.0.1) 353
infrastructure monitoring updates 352
reporting engine 349
threat center dashboard 346
Veeam Backup & Replication monitoring improvements 354
Veeam intelligence integration 356
Veeam ONE v13 architecture **360**
dual-database architecture 363
engine architecture, reporting 370
integration changes 374
network port reference 367 – 369
security changes 372
system requirements 365
v12 to v13 upgrade path 377 – 380
Veeam analytics service 361
web client and desktop client 371, 372
Veeam ONE v13 architecture, engine architecture report
data flow 371
export format changes 370
Veeam ONE v13 architecture, integration changes
Coveware recon scanner 375
intelligent diagnostics signature updates 376
Veeam intelligence 374
Veeam ONE v13 architecture, security changes

Technet24

NTLM deprecation 373
protocol modernization 372, 373
RBAC and role changes 373
REST API authentication 374
SAML-based single sign-on 373

Veeam ONE v13 architecture, system requirements
hardware support 366
monitored platform support 367
operating system 365
software prerequisites 366

Veeam Readiness program
reference link 36

Veeam Restore Operator 70

Veeam Security Best Practices Guide 195

Veeam Software Appliance (VSA) 48, 182, 207, 302
versus JeOS 88

Veeam Tape Operator 71

Veeam Updater service 75

Veeam analytics service 361
architecture 361, 362
deployment and push installation 362, 363

Veeam backup administrator role 179

Veeam backup operator role 179

Veeam cloud connect, for DRaaS 285
design considerations 285, 286

Veeam console 123
advanced tips and tricks 141
Backup Infrastructure view 134, 135
backups information 147
best practices, for backup consoles 141, 142
CTRL key 147
deployment 121, 122
Files view 136, 137
home section 124
Inventory view 133, 134
Refresh button 148, 149
secret CDP bottleneck information 146
session statistics 142 – 144
text editing 144, 145

Veeam continuous data protection (CDP)
CDP policy 252
core components 251
data flow sequence 252, 253
I/O filter cross-compatibility 280
I/O filter deployment 253 – 256
I/O journal 252
network resilience 250
pre-configuration requirements 253
prerequisites 253
proxy sizing considerations 258 – 261
source CDP proxy 251
target CDP proxy 252
VMware I/O filter 251

Veeam deployment kit 293

Veeam explorers 425
for Microsoft Active Directory 425
for Microsoft Exchange 425
for Microsoft SharePoint 425
for Microsoft SQL server 425
for Oracle 426
for PostgreSQL 426
launching 426

Veeam intelligence integration 356, 381
AI-powered solution 386
data resilience summary (v13.0.1) 357 – 360, 388
deep-data analysis agent (v13.0.1) 357
interactive dashboards and widgets 384
malware threat analysis agent (v13.0.1) 357
reporting engine 388
threat center integration and security reporting 391
VBR console integration 392
version 13.0.1 additions 393, 394
web client experience, redesigning 381

Veeam intelligent diagnostics (VID) 356, 376

Veeam restore operator role 179

Veeam universal license (VUL) 248

Veeam updater 306
approval workflow 309
deployment update, scheduling 309
enabling 307 – 309
troubleshooting scenarios 311
update history 310
update notification 309
update process, for VSA deployment 310

Veeam webpages
reference link 20

Virtual Appliance mode 21

Virtual Infrastructure 345

vSphere
limitations 244

vStorage APIs **22**
veeamhasvc.service **209**
virtual cluster IP **209**
virtual disk
restoring 422 – 425
virtual disk restore **415, 422**
volume-level restore **415, 419 – 421**

W

WAL file replication **209**
Web UI
versus Windows UI 138 – 140
Windows Backup Repositories **28**
Windows workloads **400**
Windows, moving to VSA
benefits 58
web client and desktop client **371, 372**
web client experience
appearance and theming 382
credentials manager and data collection 383
dashboard catalog redesign 381, 382
global search 383
navigation structure 383
redesigning 381
web-based user interface **68**
benefits 69, 70
capabilities 68, 69
features 68, 69
workgroup
advantages 186
disadvantages 186, 187
workgroup deployment **197, 198**

X

XFS configuration
requirements 115

Z

zero trust architecture **62**

Technet24