DEFINE • **DESIGN** • DEPLOY

FortiAnalyzer

Architecture Guide

Version 7.2

FORTINET DOCUMENT LIBRARY
https://docs.fortinet.com

FORTINET VIDEO GUIDE
https://video.fortinet.com

FORTINET BLOG
https://blog.fortinet.com

CUSTOMER SERVICE & SUPPORT
https://support.fortinet.com

FORTINET TRAINING & CERTIFICATION PROGRAM
https://www.fortinet.com/training-certification

NSE INSTITUTE
https://training.fortinet.com

FORTIGUARD CENTER
https://www.fortiguard.com

END USER LICENSE AGREEMENT
https://www.fortinet.com/doc/legal/EULA.pdf

FEEDBACK
Email: techdoc@fortinet.com

February 23, 2023
FortiAnalyzer 7.2 Architecture Guide
00-72-853410-20230223

TABLE OF CONTENTS

Change Log

Date	Change Description
2023-02-23	Initial draft.

Introduction

Executive summary

This document aims at providing design guidelines for stable and efficient FortiAnalyzer and FortiAnalyzer BigData deployments.

Starting from sizing key performance indicators, this document helps you to build an architecture on solid foundations that allows you to scale your deployment as your needs evolve over time.

Solutions and technologies

Please refer to the product datasheets for further information about FortiAnalyzer and FortiAnalyzer BigData:

- FortiAnalyzer
- FortiAnalyzer BigData

Intended audience

This guide has primarily been created for a technical audience, including system architects and design engineers who want to deploy FortiAnalyzer on solid foundations.

It assumes the reader is familiar with the basic concepts of networking, security, and FortiAnalyzer key concepts. For more information about FortiAnalyzer key concepts, see the FortiAnalyzer Administration Guide.

About this guide

This guide presents one of many possible ways to design the solution from a log rate (LPS) and storage requirements standpoint. It may omit specific steps, such as log transport, troubleshooting/health assessments, or minimum system requirements where readers must make design decisions to further configure their devices. It is recommended that readers also review supplementary material in product administration guides, example guides, cookbooks, release notes, and other documents where appropriate on the Fortinet Document Library.

The reader is required to have knowledge of corporate policies on log retention and analysis because they will be used to drive the sizing exercise.

Sizing

Sizing is a paramount step of any design exercise. Two key metrics are evaluated:

- Log rate on page 6
- Storage requirements on page 7

Log rate

Logs per second (LPS): Average number of logs per second generated in a 24-hour period that a FortiAnalyzer unit will have to sustain.

The FortiAnalyzer datasheet and FortiAnalyzer BigData datasheet provide the maximum constant log message rate that each FortiAnalyzer platform can maintain for minimum 48 hours without system performance degradation.

For existing deployments, LPS can be obtained by querying devices already deployed. For new deployments (greenfield), LPS can be estimated from either sessions/sec rate or number of users per site as described below.

To estimate LPS from sessions/sec rate:

Generally, the amount of traffic logs per second is equal to the amount of sessions per second.

If additional security features are enabled, the logs generated from each feature must be added to the total according to the table below:

Log Type	% Traffic log
Antivirus	5%
IPS	5%
Application Control	20%
Web Filtering	20%
DNS	5%

Example:

Site A generates 1500 sessions/sec, and it has Antivirus, IPS, and Application Control features enabled.

Traffic log/sec = Sessions/sec

Estimated LPS:

- Traffic (1500) + Antivirus% (75) + IPS% (75) + Application Control% (300) = Total logs/sec (1950)

To estimate LPS from number of users:

The LPS can be obtained from:

- Total number of users per site
- % of active users per day (use 50% as baseline)

Each user generates an average of 0.66 traffic logs/sec, and security features enabled must be added to the total according to the table below:

Log Type	% Traffic log
Antivirus	5%
IPS	5%
Application Control	20%
Web Filtering	20%
DNS	5%

Example:

Site A has 100 users in total, and 50% are active per day. The following security features are enabled: Antivirus, Web Filtering, and DNS.

Estimated % log per user:

- Traffic log (0.66) + Antivirus% (0.033) + Web Filtering% (0.132) + DNS% (0.033) = 0.858

Estimated LPS:

- Total users (100) * % active users (0.5) * Estimated % log per user (0.858) = 42.9

Important notes

- LPS estimated with either sessions/sec or number of users is a gross estimate in order to choose the most accurate platform for your needs. Real numbers may differ; therefore, trends should be monitored when in production.
- Specific functionality, such as SD-WAN, can increase the overall log rate by 5-10% based on the logging and monitoring configuration in place.
- A projected log volume in one to three years must be taken into account.

Storage requirements

Storage requirements: The total storage needed is directly related to the previously estimated LPS and to corporate policies on log retention and analysis.

Conventional FortiAnalyzer

In a conventional FortiAnalyzer, logs are stored in two different formats:

- Archive logs: offline logs used for log retention only
- Analytic logs: online logs indexed in SQL database and available for analytics support (FortiView, Reports, etc.)

The following table provides the average log size per log storage format (since FortiOS v6.4.3):

Format	Size
Archive Log	80 bytes
Analytic Log – Uncompressed *	600 bytes
Analytic Log – Compressed*	150 bytes

*By default, analytic logs older than 7 days are compressed. This makes them slower to retrieve, read, and display, but much more storage efficient.

The following CLI command can be used to specify after how many days the analytic logs are compressed:

```
#config system sql
   set compress-table-min-age <days>
end
```

To estimate the total storage requirement:

The estimated storage requirement is the sum of the total archive log size and total analytic log size:

- Total storage requirement = archive log size + (analytic compressed log size + analytic uncompressed log size)

To determine the total log sizes, you can use the following formulas:

- Archive log size = Log Rate * Archive Log size * 86400 seconds * number of days
- Analytic compressed log size = Log Rate * Analytic compressed size * 86400 seconds * Analytic days compressed
- Analytic uncompressed log size = Log Rate * Analytic uncompressed size * 86400 seconds * Analytic days uncompressed

Example:

Log rate = 1500 logs/sec

Desired archive period = 365 days (1 year)

Desired analytic period = 90 days (3 months)

1. Estimate the total archive size:
   - Archive size per day = 1500 logs/sec * 80 bytes * 86400 seconds = 10.37 GB
   - Total archive for 1 year = 10.37 GB * 365 days = 3.78 TB
2. Estimate the analytics size compressed (considering 83 days compressed):
   - Analytics compressed per day = 1500 logs/sec * 150 bytes * 86400 seconds = 19.44 GB
   - Analytics compressed for 83 days = 19.44 GB * 83 days = 1.61 TB
3. Estimate the analytics size uncompressed (considering the default 7 days analytics compression):
   - Analytics uncompressed per day = 1500 logs/sec *600 bytes * 86400 seconds = 77.76 GB
   - Analytics uncompressed for 7 days = 77.76 * 7 days = 0.54 TB
4. Estimate the total analytics size (compressed and uncompressed):
   - Total analytics for 3 months = Total analytics size uncompressed (1.61 TB) + Total analytics size compressed (0.54 TB) = 2.15 TB
5. Estimate the total storage requirement:
   - Total storage requirement = Total archive (3.78 TB) + Total analytics (2.15 TB) = 5.93 TB

Note: The calculation is executed using Gigabyte and Terabyte conversion rather than Gibibyte and Tebibyte.

- 1 Gigabyte = 1000^3 bytes
- 1 Gibibyte = 1024^3 bytes
- 1 Terabyte = 1000^3 Gigabyte
- 1 Tebibyte = 1024^3 Gigabyte

FortiAnalyzer BigData

The concept of archive and analytic logs doesn't apply to FortiAnalyzer BigData. Logs are compressed, replicated, and load-balanced across the hosts and available for immediate analytics.

The average log size on a FortiAnalyzer BigData is roughly 300 bytes post replication (x3) and compression.

Total storage requirements are obtained by multiplying the daily storage amount by the number of retention days needed using the below formula:

- Daily storage amount = Log Rate * 300 bytes * 86400
- Total storage requirements = daily storage amount * days

Design considerations

FortiAnalyzer and FortiAnalyzer BigData are available in either appliance or virtual machine forms. The table below indicates how key design pillars are represented:

Pillar	Physical Appliance	Virtual Machine
Performance	Proprietary hardware architecture provides the best performance experience.	Dependent on the virtual infrastructure.
Availability	Built-in software and hardware high availability and data resiliency.	Built-in software high availability and data resiliency.
Scalability	Scaling out with the addition of more chassis.	Scaling up and scaling out.
Security	Integrity of all elements is provided by the proprietary hardware architecture.	Dependent on the virtual infrastructure.
Deployment	Physical rack space needed.	Flexible deployment.

The FortiAnalyzer datasheet and FortiAnalyzer BigData datasheet provide key specifications in terms of sustained log rate/ingestion rate for both physical appliances and virtual machines.

Architectures

The following architectures are described in this guide:

- Single tier standalone on page 11
- Single tier high availability on page 12
- Multiple tiers high availability on page 13
- Multiple tiers FortiAnalyzer fabric on page 13
- FortiAnalyzer BigData on page 14

Single tier standalone

Description

- Default mode that supports all FortiAnalyzer features.
- Cost effective and easy to manage.

Limitations

- No redundancy (single point of failure).
- Scalability and redundancy limited to single device capacity.
- Scheduled backups needed for FortiAnalyzer configuration and logs.

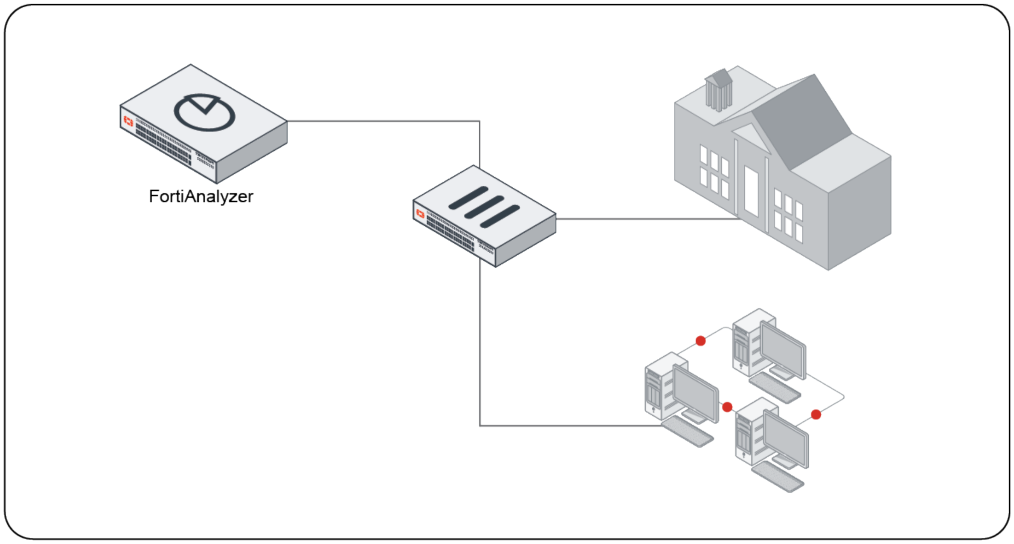

Use case

SOHO and SMB with a limited number of managed devices and logs per second (LPS) where redundancy and scalability across multiple devices is not a requirement.

Single tier high availability

Description

- Real-time redundancy for up to 4 cluster members in high availability.
- Log and data synchronized securely across cluster members.
- Alleviates the load on the primary unit by sharing resource demanding activities (such as Reports and FortiView) shared across units.
- It may be required to fulfill governmental requirements or certifications.

Limitations

- Log collection is handled by a single unit at the time.

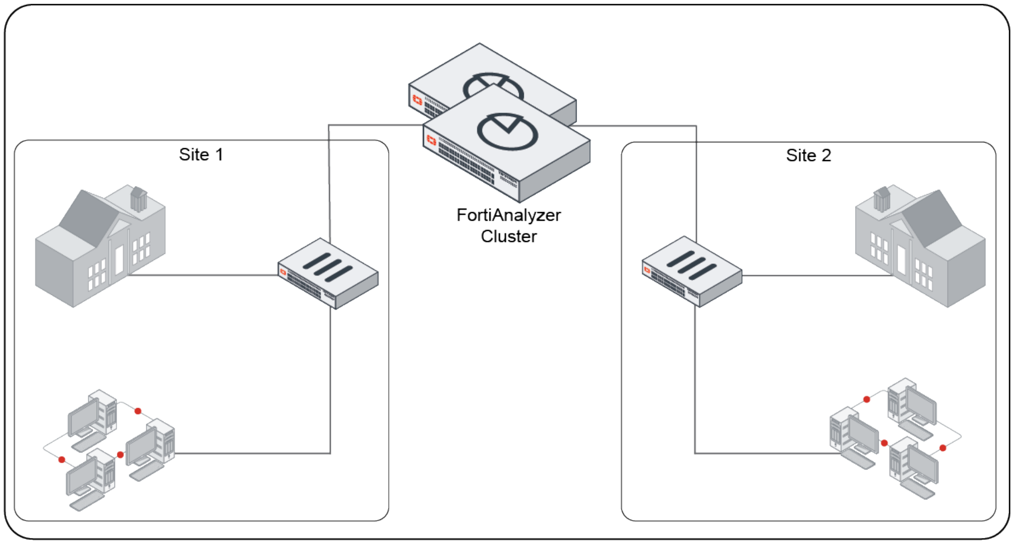

Use case

Preferred architecture for SMB/SME with a limited number of regional sites that are looking for real-time logs and data redundancy.

Multiple tiers high availability

Description

Log collection, spread across multiple layers:

- Collector layer: first layer dedicated to log collection, archiving and forwarding to the analyzer layer.
- Analyzer layer: second layer deployed in high availability and focused on analytics and reporting activities.

The FortiAnalyzer in the analyzer layer can be deployed in high availability to ensure real-time redundancy and log/data synchronization.

Limitations

- Analytics sustained log/sec rate limited to single device capacity.
- Collectors only provide historical log view (no reporting or FortiView).

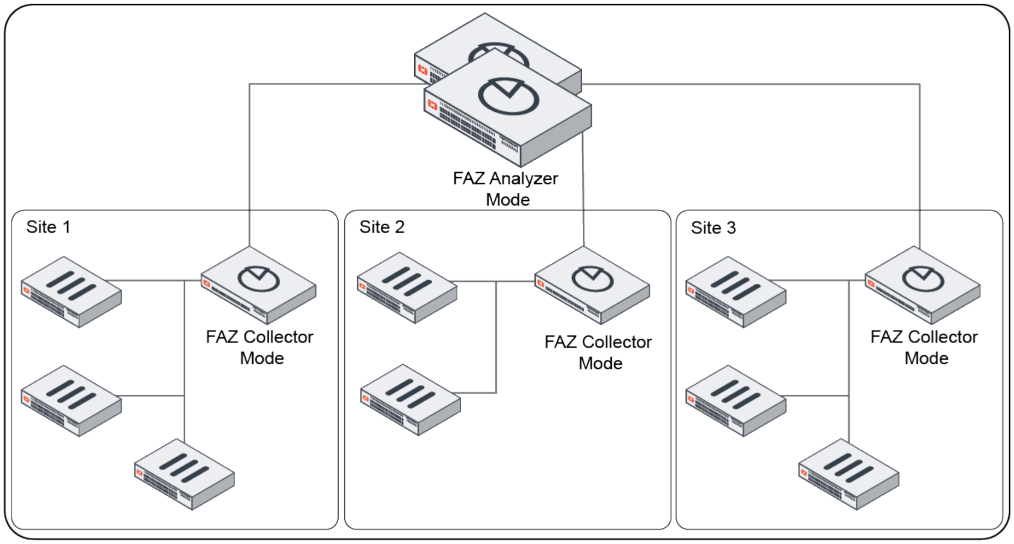

Use case

Large international companies with regional sites. The collector tier is regionally deployed to establish an OFTP connection to each managed device and to consolidate logs regionally before forwarding them to the analyzer tier. If the connectivity between the collector and analyzer tier is unavailable, logs can be buffered at the collector level.

Multiple tiers FortiAnalyzer fabric

Description

- Centralized visibility of managed devices, log view, incidents and events from a FortiAnalyzer supervisor.
- Single or multiple tiers deployment are supported (for example, Collector-Analyzer).

Limitations

- No high availability on the supervisor.
- Scalability and redundancy limited to each FortiAnalyzer deployment.
- Unable to perform configuration changes or to run automation playbooks from fabric supervisor to members.

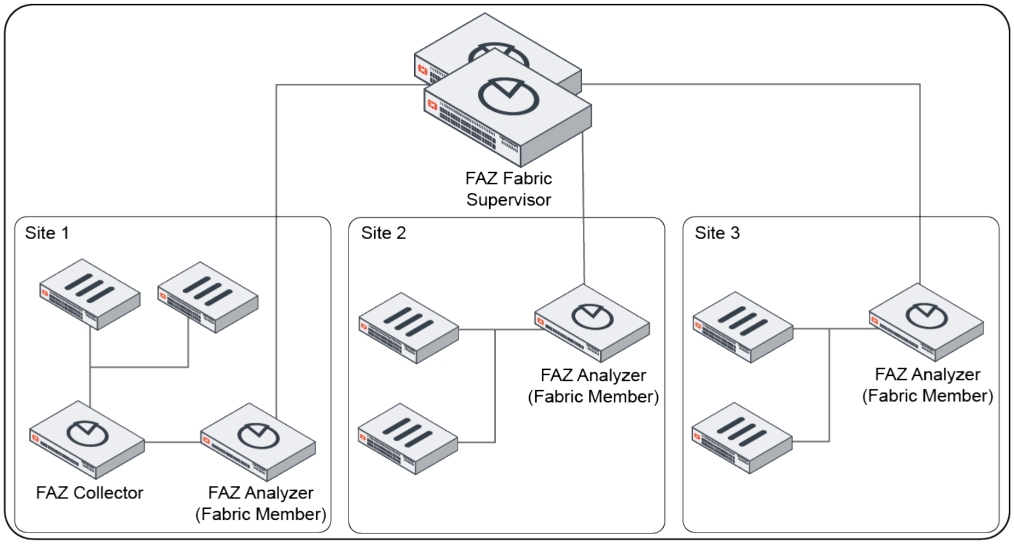
FAZ Fabric Supervisor
Site 1
Site 2
Site 3
FAZ Analyzer (Fabric Member)
FAZ Analyzer (Fabric Member)
FAZ Collector
FAZ Analyzer (Fabric Member)

Use case

A suitable architecture for multinational customers with subsidiaries worldwide. The key differentiator compared to the to the Multiple tiers high availability on page 13 architecture is that the regional SOC team can also benefit from a fully functional analyzer and not only have a historical log view available as on a collector. The collectors can be used to build points of presence in the different regions and forward the information to the central analyzer in the company's head office.

FortiAnalyzer BigData

Description

- Enterprise-grade big data software architecture (Hadoop, Spark, Kafka).
- Horizontal scalability and built-in data resiliency.
- Ideal solution for distributed enterprises and service providers with high log ingestion needs (starting at 100K logs/sec).
- Scaling out with an additional chassis allows redundancy of log collection.
- Fits both single or multiple tiers architectures; the previous examples can be used as references.

Use case

A large-scale data center and high-bandwidth deployments with high log ingestion needs, and high availability and data resiliency requirements.

Appendix A - Documentation references

FortiAnalyzer

- FortiAnalyzer document library: https://docs.fortinet.com/product/fortianalyzer
- FortiAnalyzer Private Cloud: https://docs.fortinet.com/product/fortianalyzer-private-cloud
- FortiAnalyzer Public Cloud: https://docs.fortinet.com/product/fortianalyzer-public-cloud
- FortiAnalyzer Datasheet: https://www.fortinet.com/content/dam/fortinet/assets/data-sheets/fortianalyzer.pdf

FortiAnalyzer BigData

- FortiAnalyzer BigData document library: https://docs.fortinet.com/product/fortianalyzer-bigdata
- FortiAnalyzer BigData Private Cloud: https://docs.fortinet.com/product/fortianalyzer-bigdata-private-cloud
- FortiAnalyzer BigData Datasheet: https://www.fortinet.com/content/dam/fortinet/assets/data-sheets/fortianalyzer-bd.pdf

FORTINET®

www.fortinet.com