

Administering Websense Databases

Administering Websense Databases | Web, Data, and Email Solutions | v8.0.x, v8.1.x

Websense advanced protection solutions include several databases used to store configuration information, reporting data, and other information.

This paper is designed to help database administrators understand the requirements of Websense databases across all Websense® TRITON® products.

Topics addressed in this document include:

- ◆ *Overview of Websense databases*, page 2
- ◆ *Understanding the reporting databases*, page 4
- ◆ *Reporting database specifications and recommendations*, page 8
- ◆ *How big will my reporting database be?*, page 11
- ◆ *Can I use SQL Server 2008 R2 Express?*, page 15
- ◆ *Factors that affect reporting database size*, page 17
- ◆ *Other factors that affect the performance of Websense reporting*, page 20
- ◆ *TRITON reporting database FAQs*, page 24, such as:
 - *Which database tools are required or used?*, page 24
 - *Which permissions are required?*, page 24
 - *Which database jobs are run?*, page 25
 - And more!

Overview of Websense databases

Administering Websense Databases | Web, Data, and Email Solutions | v8.0.x, v8.1.x

Applies to:	In this topic
<ul style="list-style-type: none"> ◆ TRITON AP-WEB and Web Filter & Security, v8.0.x, v8.1.x ◆ TRITON AP-DATA, v8.0.x, v8.1.x ◆ TRITON AP-EMAIL, v8.0.x, v8.1.x 	<ul style="list-style-type: none"> ◆ Websense Data Fingerprint Database, page 3

Websense TRITON solutions use a variety of databases for different purposes: configuration information, reporting data, URL categorization, fingerprinting, and forensics. Several data formats are used, including SQL, PostgreSQL, and Websense proprietary formats.

These databases include:

Database	Description
Websense reporting databases	<p>Web, Data, and Email products</p> <p>SQL Server databases that store reporting and logging data for individual TRITON Manager modules. The Data reporting database also stores configuration data. See Understanding the reporting databases, page 4.</p>
Websense TRITON Settings Database	<p>Web, Data, and Email products</p> <p>PostgreSQL database that stores global configuration and infrastructure settings that affect all TRITON Manager modules. It is installed automatically on the TRITON management server and requires no administrator configuration.</p>
Websense Master Database	<p>Web products only</p> <p>Proprietary database that contains URL categories and protocol definitions, as well as supporting information, such as risk class groupings.</p> <p>A copy of the Master Database resides on each Filtering Service machine. By default, a full update is performed daily. Incremental updates can occur much more frequently if they are enabled on the Settings > General > Database Download page in the Web module of TRITON manager.</p> <p>See “The Websense Master Database” in the Administrator Help for your Web protection solution for further details.</p>

Database	Description
Websense RTM Database	<p>Web products only</p> <p>Holds and organizes filtering data for display in Real-Time Monitor. This is an independent database (not hosted on SQL Server) installed with each RTM Client and RTM Server instance.</p> <p>Administrators can specify when Real-Time Monitor captures data on the Settings > Reporting > Preferences page in the Web module of TRITON Manager. No other aspect of database behavior is configurable.</p>
Websense Web Forensics Database	<p>TRITON AP-WEB only</p> <p>Stores details about files that may be associated with advanced malware threat activity in your network.</p> <p>Enable or disable the forensics repository, and configure its location and size on the Settings > Reporting > Dashboard page in the Web module of TRITON Manager.</p> <p>See “Configuring Dashboard reporting data” in the Administrator Help for your Web protection solution for details.</p>
Websense Data Fingerprint Database	<p>TRITON AP-DATA only</p> <p>Stores data security fingerprints.</p> <p>See Websense Data Fingerprint Database, page 3.</p>
Websense Data Forensics Database	<p>TRITON AP-DATA only</p> <p>Contains information about DLP and discovery transactions that resulted in incidents, such as the contents of an email body, including the From:, To:, and Cc: fields, as well as actual attachments. Transactions can also include Web posts, endpoint operations, and discovered as well as other events. For transactions that occurred on a Web channel, the forensics might include the URL category property. For discovery transactions, forensics might include the host name and file name.</p> <p>Configure the size and location of the forensics repository in the Data module of TRITON Manager. Navigate to the Settings > Deployment > System Modules page and click Forensics Repository under the management server.</p>

Websense Data Fingerprint Database

One of the ways that you can classify data in TRITON AP-DATA is by “fingerprinting” it. This allows the detection of sensitive information despite manipulation, reformatting, or other modification. Fingerprints enable the protection of whole or partial documents, antecedents, and derivative versions of the protected information, as well as snippets of the protected information whether cut and pasted or retyped.

When you fingerprint data, the fingerprints are stored in the Data Fingerprint Database on the TRITON management server and pushed to other Data components for fast analysis on those machines.

To tune performance, you can configure the disk space and cache size of the database on the management server in the Data module of TRITON Manager. To do so:

1. Log onto TRITON Manager and select the **Data** module.
2. Navigate to **Settings > Deployment > System Modules**.
3. Select **Primary Fingerprint Repository** under the management server.
4. Adjust the maximum disk space and cache size as needed.

Secondary repositories are maintained by Websense Content Gateway and TRITON AP-EMAIL. You can configure TRITON AP-DATA to detect fingerprints from repositories local to those machines (best practice) or from a remote machine, such as the management server.

Periodically, you must synchronize the fingerprints in the secondary repositories with the ones on the management server. How often you synchronize depends on your business needs. For best performance, select a time with low traffic volume.

To configure these settings:

1. Navigate to **Settings > Deployment > System Modules**.
2. Select **Secondary Fingerprint Repository** under the component or solution of interest.
3. Choose the repository from which to detect fingerprints.
4. Set the maximum cache size and synchronization frequency as needed.

See “Configuring the fingerprint repository” in the [TRITON AP-DATA](#) Help for instructions on configuring these settings.

Understanding the reporting databases


Administering Websense Databases | Web, Data, and Email Solutions | v8.0.x, v8.1.x

Applies to:	In this topic
<ul style="list-style-type: none"> ◆ TRITON AP-WEB and Web Filter & Security, v8.0.x, v8.1.x ◆ TRITON AP-DATA, v8.0.x, v8.1.x ◆ TRITON AP-EMAIL, v8.0.x, v8.1.x 	<ul style="list-style-type: none"> ◆ SQL Server Express on the TRITON management server, page 7 ◆ SQL Server on another machine, page 7

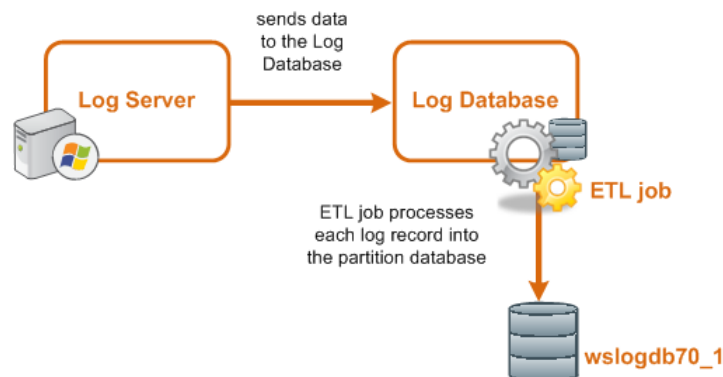
Each TRITON reporting database stores the logging data collected by a specific TRITON solution: TRITON AP-WEB, Web Filter & Security, TRITON AP-DATA, or TRITON AP-EMAIL.

- ◆ **Web Log Database** stores Internet request data collected by Log Server, such as the source, destination, time, category, risk class, action (also called disposition), bytes sent and received, and so on for use by Web reporting tools.

Log Server receives information about Internet activity from Filtering Service and initially stores it locally:



Whenever it is able, Log Server forwards the cache files to the Log Database, where the ETL job processes them into log records in a database partition:

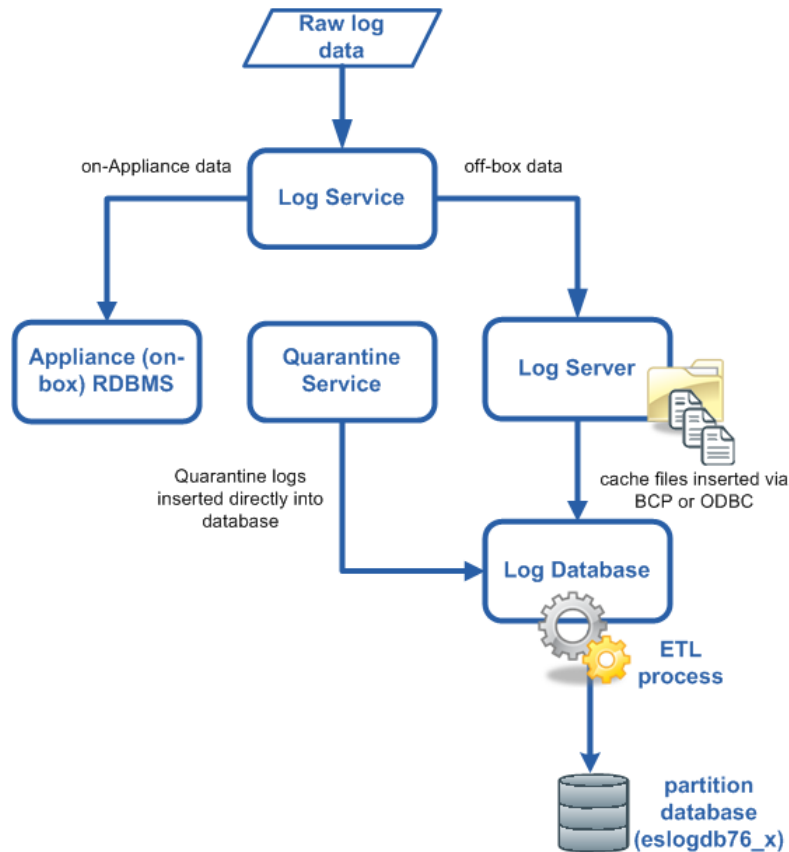


An end user who uses the Filtering Service has no direct or indirect influence over the database. Thus, although the log entry is stored in the MSSQL database, the user did not direct its storage and cannot retrieve it.

The only interface to the database itself is from the Log Server, the Reporting services, and the Manager. Filtering Service and Websense Content Gateway do not access the database, but instead send information via the Log Server.

- ◆ **Email Log Database** stores records of email traffic and the associated analysis and disposition of that traffic. TRITON AP-EMAIL reporting uses this information to generate dashboard status charts and email activity reports showing, for example, the size and volume of messages processed, message analysis results, and email source and destination.

Log and quarantine data are recorded as follows:



- ◆ **Data Incident and Configuration Database** stores information about email, Web, and other traffic that resulted in data loss prevention (DLP) policy breaches, such as the source, destination, time, status, and severity of each breach. It also stores TRITON AP-DATA policy configuration and system settings.

TRITON reporting databases are all hosted by Microsoft SQL Server. They may be hosted by the same installation and instance, or by different installations or instances.

Although Microsoft SQL 2008 R2 Express (SQL Server Express) is packaged with Websense software, enterprises with high transaction volume should purchase Microsoft SQL Server Standard or Enterprise. (See [Can I use SQL Server 2008 R2 Express?](#), page 15 for guidance.)

When you install Websense software, you can either install SQL Server Express on the TRITON management server or connect to an existing SQL Server instance on another machine.

Related topics:

- ◆ [Reporting database specifications and recommendations](#), page 8
- ◆ [How big will my reporting database be?](#), page 11
- ◆ [Can I use SQL Server 2008 R2 Express?](#), page 15
- ◆ [Factors that affect reporting database size](#), page 17

- ◆ [Other factors that affect the performance of Websense reporting, page 20](#)
- ◆ [TRITON reporting database FAQs, page 24](#)

SQL Server Express on the TRITON management server

For smaller deployments (up to 500 users), TRITON reporting databases can be installed on the TRITON management server.

The TRITON installer installs SQL Server 2008 R2 Express (32-bit) to host reporting databases when you select **Install SQL Server Express on this machine** during installation. SQL Server 2008 R2 Express is the only version of SQL Server that may reside on the TRITON management server.

If you are upgrading from a previous version of Web protection solutions and have a different version of SQL Server installed on the TRITON management server machine, move your reporting database to a different machine for improved performance and a better overall product experience.

SQL Server Express can run on the following Windows platforms. Note that Windows Server 2008 R2 and later support all TRITON products.

	Windows Server 2008 32-bit	Windows Server 2008 R2	Windows Server 2012
Data		x	x
Web	x	x	x
Email		x	x

SQL Server on another machine

In medium and larger deployments (over 500 users), the best practice is to host TRITON reporting databases on a Standard or Enterprise installation of Microsoft SQL Server. SQL Server typically resides on a dedicated machine. It may host databases for third-party applications, in addition to the Websense databases.

If SQL Server resides on a dedicated machine and there are multiple SQL instances installed, SQL Server Browser must be running.

When installing Websense software, you are prompted to provide the IP address or hostname of the SQL Server machine. You may also specify:

- ◆ An instance name
- ◆ The port to use for SQL Server communication
- ◆ Whether or not to encrypt communication with SQL Server

The following databases are supported.

	SQL Server 2008*	SQL Server 2008 R2 Express	SQL Server 2008 R2**	SQL Server 2012 SP1	SQL Server 2014***
Data	x	x	x	x	x
Web	x	x	x	x	x
Email	x	x	x	x	x

* All editions except Web, Express, and Compact; all service packs, 32- and 64-bit, excluding IA64.

**All editions except Web and Compact; all service packs, 32- and 64-bit, excluding IA64.

***SQL 2014 is supported for v8.1 only.

Clustering is supported for all supported versions of SQL Server noted in the table above. See [Can Websense reporting databases be hosted in a SQL Server cluster?](#), page 28.

Reporting database specifications and recommendations

Administering Websense Databases | Web, Data, and Email Solutions | v8.0.x, v8.1.x

Applies to:	In this topic
<ul style="list-style-type: none"> ◆ TRITON AP-WEB and Web Filter & Security, v8.0.x, v8.1.x ◆ TRITON AP-DATA, v8.0.x, v8.1.x ◆ TRITON AP-EMAIL, v8.0.x, v8.1.x 	<ul style="list-style-type: none"> ◆ Hardware specifications, page 9 ◆ RAM, page 9 ◆ Disk considerations, page 9 ◆ Virtualization, page 10

The performance of Websense reporting solutions is heavily dependent on the SQL Server, and the configuration of its underlying resources. The more you invest in the system, the better it will perform.

For optimal results, host the reporting databases on Windows Server 2012 R2 Enterprise Edition and SQL Server 2012 R2 Enterprise Edition.

Also consider the factors that follow when designing your database system.

Hardware specifications

Use hardware that meets or exceeds Microsoft's recommended (not minimum) hardware requirements for SQL Server.

Microsoft SQL Server 2014:

[https://technet.microsoft.com/en-us/library/ff928358\(v=sql.10\).aspx](https://technet.microsoft.com/en-us/library/ff928358(v=sql.10).aspx)

Microsoft SQL Server 2012:

[http://technet.microsoft.com/en-us/library/ms143506\(v=sql.110\).aspx](http://technet.microsoft.com/en-us/library/ms143506(v=sql.110).aspx)

Microsoft SQL Server 2008 R2:

[http://technet.microsoft.com/en-us/library/ms143506\(v=sql.105\).aspx](http://technet.microsoft.com/en-us/library/ms143506(v=sql.105).aspx)

Microsoft SQL Server 2008:

<http://technet.microsoft.com/en-us/library/ms143506%28SQL.100%29.aspx>

RAM

The reports that administrators generate often require that SQL Server be capable of loading, summarizing, and processing large amounts of data across multiple physical databases. Even if your organization doesn't run complex reports, if there are multiple reporting administrators, the system may frequently be asked to generate several reports concurrently. This places high demands on system memory. As a result, increasing RAM may provide noticeable performance improvements.

Keep in mind that the operating system version and/or the version of SQL Server may limit the amount of physical RAM that can be used by the system. For example, Windows Server 2008 Standard Edition supports SQL Server 2008 running with a maximum of 32 GB of RAM, no matter how much is available in hardware.

Consult your operating system and SQL Server documentation to ensure that you utilize the maximum physical RAM possible in your chosen system.

See *Other factors that affect the performance of Websense reporting*, page 20, for guidance on how the memory demands of Websense reporting may increase or decrease based on how you use it.

Disk considerations

Because database operations are often I/O-intensive, a faster, dedicated disk can improve how the database performs. Use high-performance disks whenever possible to ensure optimum database operation.

SQL Server performs better when there is minimal disk contention. For best performance, place tempdb files, reporting database files, and log files on separate physical drives with dedicated controllers. Follow SQL Server recommendations for installing and configuring SQL Server.

RAID controllers provide disk redundancy and can increase disk throughput by spreading I/O activity across multiple physical disks. If you have a high-demand system, for peak performance, use a RAID 10 configuration managed by a hardware-based RAID controller. (For systems with a lower I/O profile, RAID 5 may provide sufficient performance at a lower cost.)

For best performance:

- ◆ Use a hardware RAID controller, not software.
- ◆ Max out the RAID controller RAM cache.
- ◆ Use multi-channel links between the RAID controller and disk shelf or SAN.
- ◆ Separate the transaction logs and database files onto separate disks, arrays, or LUNs.

This means separate spindles or RAID arrays, not just separate logical drives.

If you are using a SAN, map the physical and logical LUNs to ensure you are reading and writing key components to separate physical drives. (This is to enable multiple parallel I/O operations.)

- ◆ Make sure that different database files are spread across different disks wherever possible. For very large deployments, for example, plan to map wslogdb70_1 to spindle 1, wslogdb70_2 to spindle 2, and so on.
- ◆ The disk stripe size (how files are stored across striped arrays) is also critical and is a function of both the typical read/write size and the total database size. Contact Microsoft or a sizing specialist for guidance.
- ◆ In large deployments, map different disks, arrays, or LUNs across different hardware controllers to maximize parallel operations and controller cache usage.
- ◆ If you are using spinning disks, lower seek times are better.

Virtualization

Websense products can be deployed on virtualized systems. Be aware that the use of virtualization can affect system performance. Expect up to 25% performance degradation.

For specifics on virtualization support, see the article [Virtual Machine Support](#), available from support.websense.com.

Note that Microsoft has its own support guidelines for SQL Server, which can be found here:

<http://support.microsoft.com/?id=956893>

How big will my reporting database be?

Administering Websense Databases | Web, Data, and Email Solutions | v8.0.x, v8.1.x

Applies to:	In this topic
<ul style="list-style-type: none"> ◆ TRITON AP-WEB and Web Filter & Security, v8.0.x, v8.1.x ◆ TRITON AP-DATA, v8.0.x, v8.1.x ◆ TRITON AP-EMAIL, v8.0.x, v8.1.x 	<ul style="list-style-type: none"> ◆ Web size estimates, page 11 ◆ Email size estimates, page 13 ◆ Data size estimates, page 15

When estimating the size of your TRITON reporting databases, keep the following factors in mind:

- ◆ TRITON AP-DATA logs only data violations, which typically represent a small fraction of your total web and email volume.
- ◆ TRITON AP-EMAIL logs detected spam and quarantined email messages, which typically represent up to 90% of your email volume.
- ◆ TRITON AP-WEB and Web Filter & Security log all web transactions, which typically represent about 40% of your total web traffic volume when visits (the default) are recorded rather than hits. See [Web visits, consolidation, and full URL logging](#), page 17, for details.
- ◆ If IPv6 is used in the deployment, the Web Log Database may increase in size by approximately 4%



Warning

These are averages calculated across a large number of Websense customers with widely varying deployments. While these numbers are useful for estimation, please analyze your system's actual performance after a few days or weeks to reassess your database sizing needs.

Web size estimates

The table below provides some general rules to help estimate the size of your Web Log Database, based on whether you are recording visits or hits, and URL hostnames or full URLs. By default, visits and URL hostnames are recorded.

	URL hostnames	Full URL logging
Visits	3 MB per user per month	4.5 MB per user per month
Hits	8 MB per user per month	12 MB per use per month

The size increase that accompanies full URL logging can vary widely based on the type of URL being logged. For example, the following types of traffic tend to generate very long URLs:

- ◆ Research database search and results URLs
- ◆ Search strings
- ◆ Social networking sites
- ◆ Online apps and games

Consider the example of an organization with 7500 users who wants to store data for 3 months, using monthly database partition rollover.

- ◆ Visits are enabled
- ◆ Full URL logging is enabled
- ◆ The Internet access policy is permissive
- ◆ A high percentage of Web traffic includes longer than average URLs (social networking sites and online games)

Based on this information, estimate 6 MB of data stored per user per month.

Because rollover is monthly, it is necessary to allocate enough space to hold up to 4 months of data. (This ensures that it is always possible to report on 3 months of data, even during the first days or weeks of the most recent month.)

To start, therefore, 180 GB would be needed for the database logging partitions (not including the catalog database and AMT partition).

$$6 \text{ MB} * 7500 \text{ users} * 4 \text{ months} = 180 \text{ GB}$$

In addition, calculate space for the following:

- ◆ **Database transaction logs**, whose size depends on the recovery model you choose.
 - For full recovery (not recommended), allow for the same size as the data itself (180 GB in our example).
 - For simple recovery, the log volume depends on the amount of database activity. Allow roughly 30 percent of the size of an active database partition.

In our example, the total size of 4 full partitions is 180 GB, so each partition is 45 GB in size. So for simple recovery:

$$0.3 * 45 \text{ GB} = 13.5 \text{ GB}$$

Note that this space is used only while data is being logged to the database. Inactive partitions, for example, have minimal transaction log activity. Likewise, during time periods in which little or no Internet activity is being recorded, the transaction log is very small.

- ◆ **Temporary (tempdb) storage** is used heavily during report generation and when the maintenance job is reindexing the database. The amount of tempdb space required depends heavily on the types of reports being generated.

In most environments, it is sufficient to allow twice the size of a partition:

$$45 * 2 = 90 \text{ GB}$$

- ◆ **Temporary database (tempdb) transaction logs** take up about 20 percent of the tempdb storage size. In our example:

$$0.2 * 90 \text{ GB} = 18 \text{ GB}$$

The total space required for the Web Log Database in our example, then, is:

$$180 \text{ GB} + 13.5 \text{ GB} + 90 \text{ GB} + 18 \text{ GB} = 301.5 \text{ GB}$$



Tip

As a best practice, record visits rather than hits to reduce significantly the amount of storage space required.

See [Factors that affect reporting database size, page 17](#), for an explanation of the TRITON AP-WEB visits algorithm, other data reduction options, and the full URL logging setting.

Email size estimates

For best practice, allow 1-3 MB per user per month when estimating the size of your Email Log Database.

For illustration, assume there will be 2 MB of data per user per month for 7500 users. Data will be kept for 3 months, with a scheduled monthly partition rollover.

Because rollover is monthly, it is necessary to allocate enough space to hold up to 4 months of data. (This ensures that it is always possible to report on 3 months of data, even during the first days or weeks of the most recent month.)

This means that you would need 60 GB to start:

$$2 \text{ MB} * 7500 \text{ users} * 4 \text{ months} = 60 \text{ GB}$$

In addition, you need to calculate space for the following:

- ◆ **The database transaction logs**, whose size depends on the recovery model you choose.
 - For full recovery (not recommended), allow for the same size as the data itself (60 GB in our example).
 - For simple recovery, the log volume depends on the amount of database activity. Allow roughly 30 percent of the size of the available data.

$$0.3 * 60 \text{ GB} = 18 \text{ GB}$$

- ◆ **Temporary database (tempdb) storage** is used heavily during report generation. The amount of tempdb space required depends on the reports being run. As a best practice, allow the size of the logged data for temporary storage. In our example:

$$60 \text{ GB}$$

- ◆ **Temporary database (tempdb) transaction logs** take up about 20 percent of the tempdb storage size. In our example:

$$0.2 * 60 \text{ GB} = 12 \text{ GB}$$

The total space required in our example, then, is:

$$60 \text{ GB} + 18 \text{ GB} + 60 \text{ GB} + 12 \text{ GB} = 150 \text{ GB}$$

Use the following tables to estimate storage space required for 30 days' worth of data based on email volume.

- ◆ The same requirements for partitions, logs, and temporary databases apply.
- ◆ The estimates assume an average of 5 recipients per email message.

With the Email Hybrid Module enabled for TRITON AP-EMAIL:

	10 msg/ user/day	50 msg/ user/day	100 msg/ user/day	200 msg/ user/day	400 msg/ user/day
500 users	600 MB	3300 MB	6600 MB	13300 MB	26700 MB
1000 users	1300 MB	6600 MB	13300 MB	26700 MB	53400 MB
2000 users	2600 MB	13300 MB	26700 MB	53400 MB	106800 MB
5000 users	6600 MB	33400 MB	66800 MB	133600 MB	267200 MB

Without the Email Hybrid Module:

	10 msg/ user/day	50 msg/ user/day	100 msg/ user/day	200 msg/ user/day	400 msg/ user/day
500 users	500 MB	2600 MB	5200 MB	10400 MB	20900 MB
1000 users	1000 MB	5200 MB	10400 MB	20900 MB	41800 MB
2000 users	2000 MB	10400 MB	20900 MB	41800 MB	83600 MB
5000 users	5200 MB	26100 MB	52200 MB	104500 MB	209000 MB

The calculations used are:

- ◆ With the Email Hybrid Module:

$$(0.01622 + 0.02348 * \text{RECIPIENT}) * \text{SPEED} * \text{USER}$$
- ◆ Without the Email Hybrid Module:

$$(0.0384 + 0.01322 * \text{RECIPIENT}) * \text{SPEED} * \text{USER}$$

Note that:

- ◆ USER represents the total number of users.
- ◆ RECIPIENT is the average number of recipients for each message.
- ◆ SPEED is the average message volume for each user per day.

Note that the email hybrid service drops or blocks the following types of email before they reach on-premises components:

- ◆ Email that comes from known bad (blacklisted) sources

- ◆ Email that has a very high spam score in the cloud

This significantly reduces the amount of data stored in the Email Log Database, as shown in the examples above.

Data size estimates

The Data Incident and Configuration Database is different from the Web and Email Log Databases in that it stores both configuration and log data. In addition, it logs only policy violations, not all events. In order to estimate the potential size of the Incident and Configuration Database, consider the following:

Feature	Impact on Database Size
Discovery incidents (not standard when all Websense TRITON products are installed)	14 KB per incident Typically, no more than 3.5 GB total
Network and endpoint incidents	10 KB per incident Typically, 30 KB per user per month
User directory import data	8 KB per record Typically, no more than 1 GB total
Configuration data	200 MB

As a general guideline, it is unlikely that your Data Incident and Configuration Database will require more than 9 GB of storage total, no matter how many users are in your environment.

The volume of configuration and incident data varies based on your use of TRITON AP-DATA features as explained in [Factors that affect reporting database size, page 17](#).

Can I use SQL Server 2008 R2 Express?

Administering Websense Databases | Web, Data, and Email Solutions | v8.0.x, v8.1.x

Applies to:

- ◆ TRITON AP-WEB and Web Filter & Security, v8.0.x, v8.1.x
- ◆ TRITON AP-DATA, v8.0.x, v8.1.x
- ◆ TRITON AP-EMAIL, v8.0.x, v8.1.x

Microsoft SQL Server 2008 R2 Express is bundled into the Websense TRITON Unified Installer.

Regardless of the resources available on the machine hosting the database engine, Microsoft limits the resources that SQL Server 2008 R2 Express can use to:

- ◆ 1 GB of RAM
- ◆ 1 physical CPU
- ◆ 10 GB of data per database

This results in practical limits on the amount of data that can be stored in SQL Server 2008 R2 Express, while still maintaining acceptable performance when generating reports and processing log data. For best performance, no more than 30 GB of data should be stored in SQL Server 2008 R2 Express across all Websense modules. The sizing recommendations below are based on this assumption.

For an explanation of particular demands that Websense reporting places on available memory, see [Other factors that affect the performance of Websense reporting, page 20](#).

Use the charts below to see how much data you can store using SQL Server Express based on the Websense module or modules you are using. You can then determine if that limit still meets your data retention requirements. If so, you can use SQL Server 2008 R2 Express.



Note

Standalone TRITON AP-DATA installations can support up to 5000 users with SQL Server 2008 R2 Express.

Users	Web*	Web and Data	Data and Email**	Web, Data, and Email
2000+	Do not use	Do not use	Do not use	Do not use
2000	45 days	30 days	Do not use	Do not use
1500	60 days	45 days	Do not use	Do not use
1000	90 days	60 days	Do not use	Do not use
750	4 months	3 months	Do not use	Do not use
500	6 months	4 months	30 days	Do not use
250	1 year	8 months	60 days	30 days
100	1 year+	1 year	5 months	3 months

* Web: Actual size is based on Web requests processed per day, which can be extrapolated from the peak requests per second in your network. If you are not able to determine these numbers, Websense has found that the ratio of users to their peak number of Web requests per second (as opposed to average number per second) is 10 to 1. This is an acceptable ratio to use for the purposes of estimating the number of users that this traffic represents in the average network. Note that this calculation is based on a standard curve distribution of traffic throughout the day, where the peak is mid-day and there is almost no traffic throughout the night.

** Email: Actual sizing is based on total email messages sent and received per day, including spam. Websense has found that an estimate of 1 message per user per minute (including spam) is an acceptable estimate for the purposes of estimating the email volume generated for any given number of users.

Factors that affect reporting database size

Administering Websense Databases | Web, Data, and Email Solutions | v8.0.x, v8.1.x

Applies to:	In this topic
<ul style="list-style-type: none"> ◆ TRITON AP-WEB and Web Filter & Security, v8.0.x, v8.1.x ◆ TRITON AP-DATA, v8.0.x, v8.1.x ◆ TRITON AP-EMAIL, v8.0.x, v8.1.x 	<ul style="list-style-type: none"> ◆ Web visits, consolidation, and full URL logging, page 17 ◆ Email sizing factors, page 17 ◆ Data sizing factors, page 18

Web visits, consolidation, and full URL logging

TRITON AP-WEB and Web Filter & Security use proprietary algorithms to reduce the volume of log data in order to achieve a balance between visibility into users' Web browsing activity and the size and performance of the Log Database.

- ◆ When you enable **visits**, Log Server combines the individual elements that create a Web page (such as graphics and advertisements) into a single log record that includes bandwidth information for all elements of the visit.

When this option is disabled, you instead log **hits**. In this case, a separate log record is created for each HTTP request generated to display different page elements, including graphics, advertisements, embedded videos, and so on. This creates a much larger Log Database that grows rapidly.

Disabling visits can increase the total amount of data stored in the Log Database by a factor of 2.5.

- ◆ To further reduce the size of the database, enable **log record consolidation**. This combines multiple, similar Internet requests into a single log record, reducing the granularity of reporting data.
- ◆ By default, Web products log only the URL hostname for each request, instead of the full URL. Storing the full URLs provides more visibility into where users are going within a particular website, but increases the Log Database storage demands.

Enabling full URL logging can increase the size of each record by 50%.

For information about more ways to either reduce the size of the Log Database or increase the amount of data recorded, refer to: [Log Database sizing guidance](#)

Email sizing factors

Email hybrid service

The TRITON AP-EMAIL hybrid service (included with the Email Hybrid Module) drops email that comes from known bad (blacklisted) sources and blocks email with a

very high spam score in the cloud before it ever reaches the email appliance. This reduces the amount of data stored in the Email Log Database for reporting by 30 MB per user per month.

Above average email traffic: recipients, quarantined messages, or spam

The sizing guidelines above are based on the following assumptions about the email traffic handled by TRITON AP-EMAIL. These assumptions are derived from the average email traffic pattern of Websense customers over time.

- ◆ There are an average of 5 recipients for each email message.
- ◆ When the email hybrid service is not enabled, the ratio among spam messages, infected messages, and clean messages is 85-to-1-to-14.
- ◆ The email hybrid service scans only inbound email traffic, and it can block 25 percent of spam.
- ◆ All outbound and internal email messages are clean.

Note that TRITON AP-EMAIL counts the number of recipients for each message rather than the number of messages sent. Each recipient is counted as a transaction.

If the pattern of email traffic in your organization exceeds these averages, your storage capacity will vary.

Data sizing factors

Number of discovery incidents

TRITON AP-DATA limits the number of discovery incidents that can be stored in the Data Incident and Configuration Database in order to prevent improperly configured discovery policies from flooding the database. By default this limit is set to 1 million incidents. **If you are using SQL Server Express, you should reduce this number to 250,000.**

To do this:

1. Log onto TRITON Manager.
2. Select the **Data** tab.
3. Select **Settings > General > System**.
4. Select the **Reporting** option from the System pane.
5. Select the **Discovery** tab.
6. Adjust the Maximum Discovery Incidents field.

Refer to “Setting preferences for discovery incidents” in the [TRITON AP-DATA Help](#) for more information.



Note

TRITON AP-DATA Discover is not included in combined web, data, and email TRITON solutions. It is an add-on feature that requires a separate subscription.

Rate of network and endpoint incidents

The rate of network and endpoint incidents detected varies widely across Websense customers. The sizing guidelines above are based on an average incident rate of 1 per user every 10 days (an incident is a policy violation). For best practice, periodically review the actual incident rate in the database to gauge how closely your environment matches this average, and then adjust your database storage requirements based on the actual data in your environment.

Do this by examining the Incident Trends report found in the Data module of TRITON Manager under **Main > Reporting**.



Note

Data Endpoint is not included in combined web, data, and email TRITON solutions. It is an add-on feature that requires a separate subscription.

The TRITON AP-DATA database stores data in partitions per each calendar quarter. You can have 1 active partition for the current quarter.

If you are using Microsoft SQL Server Standard or Enterprise for your TRITON database, you can have a up to 8 online partitions (approximately 2 years), but if you are using SQL Server Express, you can have only 4 (approximately 1 year). (Online partitions are partitions that can be used to show reports and log data).

For both databases, you can have up to 12 archived partitions representing 3 years of records, and 4 restored partitions (1 year).

Partition type	Microsoft SQL Server Standard or Enterprise	SQL Server Express
Active	1 partition (current quarter)	1 partition (current quarter)
Online	up to 8 partitions (2 years)	up to 4 partitions (1 year)
Restored	up to 4 partitions (1 year)	up to 4 partitions (1 year)
Archived	up to 12 partitions (3 years)	up to 12 partitions (3 years)
Total available managed partitions	25	21

Refer to “Archiving incidents” in the [TRITON AP-DATA](#) Help for more information on archiving. For instructions on setting the maximum disk space allowed for the incident archive, refer to “Configuring the incident archive”.

Size of user directory import

To support user-based policy and reporting, TRITON AP-DATA imports entries from your user directory—such as Active Directory or Domino—into the Configuration Database. Depending on the size and design of your user directory, this can result in database space being consumed by entries that are not needed by TRITON AP-DATA. To reduce the number of imported user directory entries:

- ◆ Configure TRITON AP-DATA to import entries from a more specific root context that is deeper in the tree than the directory’s root context.
- ◆ Restrict the user attributes that are imported by specifying specific attributes to import; or eliminate them altogether by disabling the import of user attributes.

To configure user directory settings:

1. Log onto TRITON Manager.
2. Select the **Data** tab.
3. Select **Settings > General > System**.
4. Select the **User Directories** option from the System pane.
5. Select the user directory to edit.
6. Modify or add a root naming context.
7. Modify the user attributes settings.

Refer to the “Adding a new user directory server” section in the [TRITON AP-DATA](#) Help for information on configuring these settings.

Other factors that affect the performance of Websense reporting

Administering Websense Databases | Web, Data, and Email Solutions | v8.0.x, v8.1.x

Applies to:	In this topic
◆ TRITON AP-WEB and Web Filter & Security, v8.0.x, v8.1.x	◆ <i>Factors affecting Web reporting performance, page 21</i>
◆ TRITON AP-DATA, v8.0.x, v8.1.x	◆ <i>Factors affecting Email reporting performance, page 22</i>
◆ TRITON AP-EMAIL, v8.0.x, v8.1.x	◆ <i>Factors affecting Data reporting performance, page 23</i>

Factors affecting Web reporting performance

Users' Web browsing behavior

Web browsing behavior varies widely among Websense customers. Periodically review your database performance, your reporting needs, and the actual data in the database so you can identify ways to reduce the demands on your reporting system.

Selective logging

TRITON AP-WEB and Web Filter & Security allow you to reduce the demands on your reporting system by not logging traffic to websites in selected categories. For example, online retailers might disable logging for Shopping categories. This can result in a large reduction in the amount of data that has to be stored and managed by Web reporting.

For information on configuring selective logging, refer to “Configuring how requests are logged” in the [Administrator Help](#) for your Web protection solution.

Use of Detail Reports over long time periods

Reports of Web browsing activity over long time periods (weeks and/or months) require much more memory, processor time, and disk I/O to generate. For better performance, run summary reports across long time periods on a regular schedule, then use detail reports only for investigating specific users in shorter time periods. If you have business requirements that demand generating detail reports across a large time window, you can:

- ◆ Schedule the reports to run during low-activity periods in your network.
- ◆ Invest in more hardware resources for your reporting system.

Number of scheduled reports or number of delegated reporting administrators

If you have several delegated administrators that use reporting each day or create several scheduled reports to run each night, this can degrade the performance of your Websense reporting tools. If you meet these usage profiles, consider investing in more hardware resources for your reporting system: more RAM, faster disks, faster CPUs, and higher-end versions of SQL Server and Windows that support more hardware.

Geographical location of users

If you have users distributed among multiple physical locations and your business does not require unified reporting across all users, consider deploying separate Log Server and Log Database instances in each location.

Calculation of Internet browse time

By default, a database job calculates Internet browse time at 2 a.m. for the previous day's activity. This is a memory-, processor-, and disk I/O-intensive activity. If you

don't use Internet browse time to manage your users' Web browsing activity, consider disabling Internet browse time to improve the performance of your reporting system.

For information on configuring Internet browse time, refer to "Configuring Internet browse time options" in the [Administrator Help](#) for your Web protection solution.

Partitioning

The Log Database is segmented into partitions for easier data management. Depending on the time period covered by a report, Websense may need to query multiple partitions. This may make report generation less efficient.

By default, a new Log Database partition is created when the current partition size reaches 5 GB (3GB if you use SQL Server Express). (With Standard and Enterprise versions of SQL Server, you can also configure the Log Database to roll over at a specific time interval.)

Review the size and content of the partitions in the database after your system has been installed and receiving data for a few days, then tune the partitioning configuration (rollover size or time period) accordingly.

For information on managing partitions, refer to "Configuring database partition options" in the [Administrator Help](#) for your Web protection solution.



Note

Due to Microsoft SQL Server restrictions, you are limited to 70 active log database partitions. You may retain more than 70 partitions, but only 70 may be active (enabled) at any one time.

Hybrid service users

The sizing guidelines in this document include logs generated by users managed by the TRITON AP-WEB hybrid service (Web Hybrid module). When sizing your reporting system, do not forget to include those users.

Configuration options that affect Log Database sizing, including selective logging, logging visits, and full URL logging, also apply to hybrid log records, so no special consideration needs to be made for those users.

Factors affecting Email reporting performance

Number of scheduled or custom presentation reports

If you create a large number of scheduled reports to run each night (more than 10) or use a large number of custom presentation reports (more than 10) each day, this can affect reporting performance. In particular, the following 3 reports place high demands on system resources:

- ◆ Top *n* External Recipients by Message Volume

- ◆ Top n External Recipients by Message Size
- ◆ Top n Data Loss Prevention Violations by Volume

If you meet these usage profiles, consider investing in more hardware resources for your reporting system: more RAM, faster disks, faster CPUs, and higher-end versions of SQL Server and Windows that support more hardware.

Partitioning

The Email Log Database is segmented into partitions for easier data management. Depending on the time period covered by a report, Websense reporting tools may need to query multiple partitions. Running such reports may be inefficient.

By default, Websense creates a new partition within the Email Log Database after storing 5 GB of data in a single partition. You can also configure TRITON AP-EMAIL to create a new partition based on a weekly or monthly time interval.

Review the size and content of the partitions in the database after your system has been installed and receiving data for a few days, then tune the partitioning configuration accordingly.

For information on managing partitions in TRITON AP-EMAIL, refer to “Configuring Log Database Options” in the [TRITON AP-EMAIL](#) Help.



Note

Due to Microsoft SQL Server restrictions, you are limited to 70 active log database partitions. You may retain more than 70 partitions, but only 70 may be active (enabled) at any one time.

Factors affecting Data reporting performance

TRITON AP-DATA automatically archives database partitions containing older data to reduce storage requirements and maintain a high performance reporting experience. To further reduce the data storage requirements, you can choose to create archive partitions sooner and keep fewer concurrent restored archives.

Refer to “Archiving incidents” in the [TRITON AP-DATA](#) Help for more information on archiving.

TRITON reporting database FAQs

Administering Websense Databases | Web, Data, and Email Solutions | v8.0.x, v8.1.x

Applies to:	In this topic
<ul style="list-style-type: none"> ◆ TRITON AP-WEB and Web Filter & Security, v8.0.x, v8.1.x ◆ TRITON AP-DATA, v8.0.x, v8.1.x ◆ TRITON AP-EMAIL, v8.0.x, v8.1.x 	<ul style="list-style-type: none"> ◆ Which database tools are required or used?, page 24 ◆ Which permissions are required?, page 24 ◆ Which database jobs are run?, page 25 ◆ How does the installer set up each database?, page 26 ◆ How big should the database partitions be?, page 27 ◆ How many partitions can be accessed at the same time?, page 27 ◆ How do I configure partition rollover?, page 27 ◆ What if I need more partitions to run reports?, page 28 ◆ Do Websense databases use named instances?, page 28 ◆ Can Websense reporting databases be hosted in a SQL Server cluster?, page 28

Which database tools are required or used?

Websense reporting components connect to the SQL Server database engine as clients and perform standard Transact-SQL commands and stored procedures.

TRITON AP-WEB and TRITON AP-EMAIL may use 2 database utilities:

- ◆ **bcp** to use bulk insertion for adding logs to the database.
- ◆ **osql** to run SQL scripts during Log Database installation.

Which permissions are required?

During TRITON AP-DATA installation, modification, or repair, the account used for database creation and access needs **sysadmin** server role membership. Also, **Backup database** permission on the **master** database is required for installation only. After installation, the account privileges can be reduced to the **db_owner** of the newly

created databases, and no access to any other user database except system databases such as master, tempdb, and model is required. Additionally the **dbcreator** server role should be granted to enable backup and restore functionality.

If you're using SQL Server 2012 Standard or Enterprise, 2008, or 2008 R2 to install the Web Log Server and Email Log Server, the user account that owns the Websense database must:

- ◆ Be a member of the **dbcreator** server role
- ◆ In the **msdb** database:
 - Have membership in the **db_datareader** role
 - Have membership in one of the following roles:
 - SQLAgentUser Role
 - SQLAgentReader Role
 - SQLAgentOperator Role

For SQL Server 2008 R2 Express, the user account requires the **sysadmin** server role.

Which database jobs are run?

The following database jobs are installed with the Web Log Database and Email Log Database:

- ◆ The Extract, Transform, and Load (ETL) job runs continually, receiving data from Log Server, processing it, and then inserting it into the partition database. The ETL job must be running to process log records into the Log Database.
- ◆ The database maintenance job performs database maintenance tasks. This job runs nightly, by default.

ETL jobs are run, then re-run 10 seconds after they finish for SQL Server Standard and Enterprise. For SQL Server Express, 60 seconds elapse between completion of one job and start of the next.

Maintenance jobs are run once every night by default. The jobs are run automatically.

The Web Log Database also installs the following jobs:

- ◆ The Internet browse time (IBT) job analyzes data and calculates browse time for each user. The IBT database job is resource intensive, requiring significant server resources. This job runs nightly.
- ◆ When trend data retention is enabled, the trend job uses daily trend data created by the ETL job to update weekly, monthly, and yearly trend records for use in presentation reports. This job runs nightly.

Even when trend data retention is disabled, the trend job processes data from the threats (AMT) partition to provide trend data on the Threats dashboard.

- ◆ The Advanced Malware Threat (AMT) ETL job receives, processes, and inserts data into the threats partition database. Only log records that include a severity ranking are recorded in the threats partition. Data from this partition is used to populate the Threats dashboard.

When configuring the start time for the (Web and Email) maintenance job and the (Web) Internet browse time job, consider system resources, and IT maintenance tasks and their duration. These jobs can be resource intensive and time consuming, so they can have a negative impact on logging and reporting performance.

Both Log Databases require either the SQL Server Agent service (SQL Server Standard or Enterprise) or Service Broker (SQL Server Express) to run database jobs.

How does the installer set up each database?

The TRITON reporting databases should allow TCP and trusted-mode connections from the TRITON management server, Email Log Server, and Web Log Server, as well as from the any email-capable V-Series appliance (Email or Web and Email mode).

Web Log Database

By default, the Web Log Database includes one catalog database, one standard logging partition database, and one threats (AMT) partition database. Typically, multiple standard logging partition databases are created as Internet activity is recorded.

- ◆ The catalog database provides a single connection point for the various components that need to access the Log Database: Log Server, presentation reports, and investigative reports configuration. It also contains definitions for the following:
 - Category names
 - Risk classes
 - Users
 - User-to-group mapping
 - Database job information

The catalog database also maintains a list of all the database partitions.

- ◆ Standard logging partitions store the individual log records of Internet activity. New partitions are created based on size (3 or 5 GB, by default) or date interval.
- ◆ The threats (AMT) partition stores information about requests that have been assigned a severity level, and is used to populate the Threats dashboard.

Email Log Database

The Email Log Database includes one catalog database and (initially) a standard logging partition.

- ◆ The catalog database provides a single connection point for the various components that need to access the Log Database: Log Server, the TRITON AP-EMAIL quarantine daemon, and the Email module of TRITON Manager (presentation reports, dashboard, log database configuration page). It also includes definitions for the following:
 - TRITON AP-EMAIL actions

- Mail direction
- Message type
- DLP severity level
- Email appliance mapping
- Email policies
- Rules
- Viruses
- DLP policy names
- IP addresses
- Email addresses
- Domains
- Database jobs

The catalog database also maintains a list of all the database partitions.

- ◆ Database partitions store the individual log records, including connection log, message log, policy log, delivery log, status log, and hybrid service status log. New partitions are created based on size (5 GB, by default) or date interval.

How big should the database partitions be?

For Web, see [Partitioning, page 22](#).

For Email, see [Partitioning, page 23](#).

For Data, see [Rate of network and endpoint incidents, page 19](#).

How many partitions can be accessed at the same time?

TRITON AP-DATA maintains incident partitions independently of the database engine, based on quarters (3-month periods). By default, SQL Server Express maintains 8 partitions are online simultaneously, and other SQL Server editions maintain 12 partitions online. You can choose to move any number of partitions online simultaneously as long as your disk space and SQL Server database permit it.

With Web and Email solutions, you can access all enabled partitions.

How do I configure partition rollover?

With Web and Email solutions, partition rollover can occur automatically when partitions reach a specified size or (SQL Server Standard or Enterprise) date.

- ◆ When partition rollover is based on size, all log records are inserted into the most recent active partition that satisfies the size rule. When the partition reaches the designated maximum size, a new partition is created.
- ◆ When partition rollover is based on date, new partitions are created according to the established cycle. For example, if the rollover option is monthly, a new

partition is created as soon as any records are received for the new month. Log records are inserted into the appropriate partition based on date.

Partition rollover can also be initiated manually.

For information about configuring automatic or manual rollover, see:

- ◆ “Configuring database partition options” in the [Administrator Help](#) for your Web protection solution.
- ◆ “Configuring Log Database options” in the [TRITON AP-EMAIL Help](#).

For Data solutions, partition rollover is configured on the **Settings > Archive** page in the Data module of TRITON Manager. Here, you configure when to create an archive partition and when to restore it. For instructions, refer to “Archiving incidents” in the [TRITON AP-DATA Help](#).

What if I need more partitions to run reports?

For Web and Email solutions, the available Log Database partitions, both enabled and disabled, are listed on the **Settings > Reporting > Log Database** page in the respective Web and Email modules of TRITON Manager. To include data from a disabled partition, first enable it, then run the report. You can use this page to disable the partition again once you have retrieved the desired data.

For TRITON AP-DATA, when you want to run a report and some or all of the data you want is stored in an offline partition, you must bring that partition online, or the generated report will not contain all the data you need.

Do Websense databases use named instances?

If you are using SQL Server Standard or Enterprise to host your Websense reporting databases:

- ◆ The Web Log Database can be hosted by one instance, and the Email Log Database and Data Incident and Configuration Database can be hosted elsewhere (on another server or instance).
- ◆ All 3 reporting databases can be hosted by the same SQL Server instance.

Can Websense reporting databases be hosted in a SQL Server cluster?

If your organization uses a SQL Server cluster to provide failover for your database servers, the Websense reporting databases can be hosted by the cluster if:

- ◆ A virtual IP address is assigned to the cluster.
- ◆ The data managed by SQL Server is housed on a reliable shared disk array.

When you install reporting components in a network that uses a SQL Server cluster, it is imperative that the cluster’s **virtual IP address** is used to configure the reporting

database connection. When this is done, reporting data is sent to SQL Server via the virtual IP address.

If you configure Websense reporting components (like Web and Email Log Server) to use the IP address of an individual node in the cluster, they cannot take advantage of the failover protection of the cluster.

- ◆ If the configured node becomes unavailable, reporting components cannot process data into the reporting database.
- ◆ For Web and Email solutions, log cache files continue to be saved on the Log Server machine. These files can build up quickly and fill the disk, causing additional problems.

When failover occurs, reporting components must wait briefly while the secondary SQL Server is made primary. When SQL Server begins accepting data over the virtual IP address again, reporting data is once again sent successfully.

This pause in recording data occurs both when failover occurs in a SQL Server cluster and when a standalone SQL Server installation fails and is later brought back online. Any records that were **actively being processed** into the reporting database when the primary SQL Server fails are lost.

- ◆ For Web and Email solutions, if you are using BCP log insertion, the loss is generally a full batch (log cache file) of filtering records.
- ◆ With ODBC log insertion, fewer records may be lost.

