

Sun Blogs: A Sun Java™ System Web Server 7.0 Reference Deployment

October 2008

Abstract

This article presents an overview of the architecture and configuration of the Sun Blogs web site, a blogging service dedicated to Sun employees worldwide. It also provides usage statistics and information about the system administrator's experience managing one of the busiest externally facing Web 2.0 site deployments at Sun.



Table of Contents

Introduction.....	1
Sun Blogs Statistics.....	1
System Architecture	3
Sun Blogs Configuration.....	4
Sun Blogs Site Administrator's Perspective.....	6
Conclusion.....	7
About the Authors.....	7
For More Information.....	7
Appendix	9
Licensing Information and Third-Party Trademarks.....	11

Executive Summary

Sun Blogs, available at <http://blogs.sun.com>, is a blogging service dedicated to Sun employees worldwide, hosted and managed by sun.com engineering. Blogs.sun.com is powered by Sun Java™ System Web Server 7.0 Update 2 (hereafter, Web Server) running on a pair of Sun Fire™ T2000 servers with UltraSPARC® T2 processors. The site also uses a custom version of the open source blogging server Apache Roller Weblogger Version 4.0.0.12 (<http://rollerweblogger.org/project/>), MySQL™ database server, and Memcached. The service is integrated with Sun's secure single sign-on identity system backed by a worldwide employee LDAP directory for authentication and authorization.

Introduction

This article presents an overview of the Sun Blogs site's usage statistics, systems architecture, and Web Server configuration. It also provides information about the system administrator's experience managing one of the busiest externally facing Web 2.0 site deployments at Sun.

Sun Blogs Statistics

Overall, as of September 2008, Sun Blogs hosted the following activity.

Table 1. Sun Blogs Statistics

Total blogs:	4770
Total bloggers:	6370
Total entries:	114689
Total comments:	121443

The graph in Figure 1 and Table 2 show monthly raw access log statistics, analyzed by The Webalizer (<http://www.webalizer.com/>). For example, in August 2008, the Sun Blogs site (blogs.sun.com) served 125 million hits, averaging roughly 4 million per day.

Usage Statistics – blogs.sun.com

Summary Period: Last 12 Months

Generated 25-Sep-2008 04:05 PDT

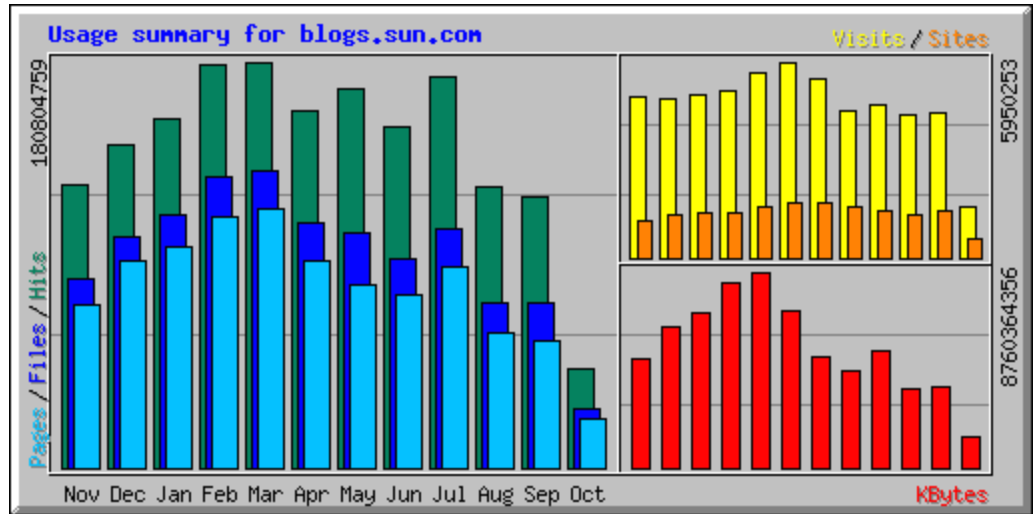


Figure 1. Usage Graph

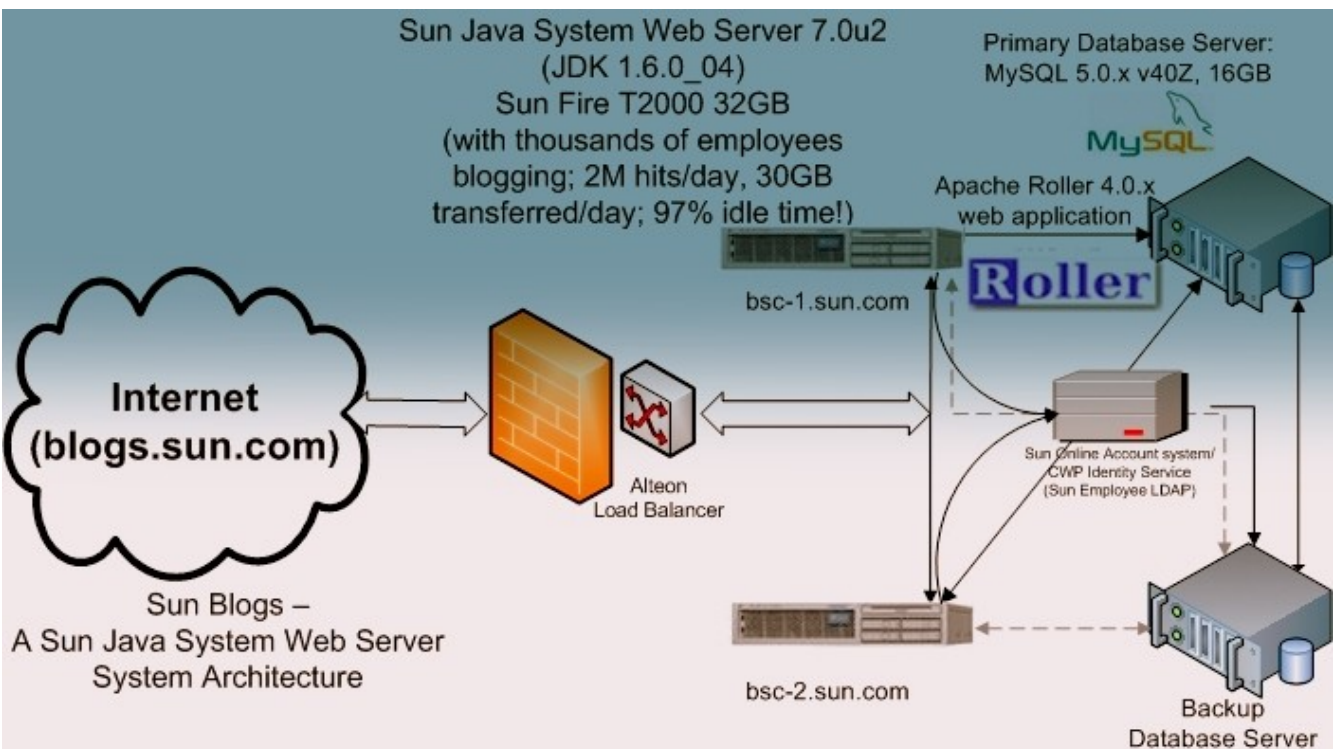
Table 2. Summary by Month

Summary by Month										
Month	Daily Avg				Monthly Totals					
	Hits	Files	Pages	Visits	Sites	Kbytes	Visits	Pages	Files	Hits
Sep 2008	4022310	2443148	1883788	147889	1168757	2909315631	3549350	45210923	58635557	96535440
Aug 2008	4047952	2361057	1935620	140051	1305519	3513460807	4341590	60004224	73192789	125486513
Jul 2008	5604441	3440767	2898110	148960	1396640	5198567170	4617788	89841411	106663792	173737681
Jun 2008	5052969	3105905	2569418	149685	1528880	4376355198	4490571	77082567	93177159	151589087
May 2008	5433226	3385538	2630128	175367	1681645	4969279215	5436397	81533978	104951686	168430031
Apr 2008	5313183	3637188	3082513	198341	1663764	7015416978	5950253	92475392	109115657	159395490
Mar 2008	5832411	4261262	3722064	181431	1519196	8760364356	5624382	115384008	132099130	180804759
Feb 2008	6194570	4463828	3856512	174737	1361481	8268537824	5067384	111838874	129451025	179642549
Jan 2008	5011329	3627057	3179296	160103	1356705	6895886673	4963215	98558187	112438772	155351214
Dec 2007	4649715	3318316	2959920	155366	1302458	6278111184	4816346	91757541	102867823	144141173
Nov 2007	4214918	2812264	2414724	162161	1145694	4892089048	4864845	72441748	84367928	126447556
Oct 2007	3238436	1915052	1421577	150558	1175550	2471445875	4667313	44068891	59366641	100391529
Totals						65548829959	58389434	980197744	1166327959	1761953022

System Architecture

Figure 2 shows the system architecture behind the Sun Blogs service (a mature instantiation of the Apache Roller Weblogger code). It is deployed on a pair of Sun Fire T2000 servers with 32 Gbytes of RAM, negligible local disk, and a shared file system (NFS from Sun Cluster), running Web Server 7.0 Update 2. The service is behind a hardware load balancer splitting traffic evenly. A ton of network gear on the front, from Juniper Network Netscreen firewalls to Nortel Alteon load balancers, does a lot of the heavy network work; the gear is fully redundant all the way through from the provider to the actual machines. The Sun Blogs site plans to switch soon to Sun Fire X4450 servers to take advantage of higher single-thread performance, because the Sun Fire T2000 servers are fairly under-utilized.

The database that powers all this is MySQL on Sun Fire V40z servers with 16 Gbytes of RAM and lots of internal disk. Memcached is used for optimal performance. No external storage is used. The database servers are paired (there are two of them) running in hot/standby mode.



Note: V40z stands for Sun Fire V40z server

Figure 2. Usage Graph

Sun Blogs Configuration

The `server.xml` configuration file contains information related to server runtime behavior, such as number of threads, number of concurrent connections, and Java Virtual Machine (JVM™) tuning for the built-in web container to serve servlet and JSP™ applications. This section shows the configuration changes that the Sun Blogs team needed to make to get the desired site performance.

Here are relevant sections of the Web Server `server.xml` configuration file, with inline comments. For the sake of brevity, only selected attributes are shown.

```
<server>
  <log>
    <log-file>../logs/errors</log-file>
    <archive-suffix>.%Y%m%d</archive-suffix>
    <log-level>info</log-level>
  </log>

  <http>
    <output-buffer-size>40960</output-buffer-size>
  </http>

  <!-- Typical worker thread pool with higher stack size not uncommon for
Java web applications with deeper invocation stack depths -->
  <thread-pool>
    <max-threads>250</max-threads>
    <stack-size>256000</stack-size>
  </thread-pool>

  <!-- Java VM settings -->
  <jvm>
    <java-home>../jdk</java-home>
    <server-class-path>...</server-class-path>
    <debug>>false</debug>
    <debug-jvm-options>-Xdebug -Xrunjdpw:
transport=dt_socket,server=y,suspend=n,address=7896</debug-jvm-options>
    <jvm-options>-java.security.auth.login.config=login.conf</jvm-options>

    <!-- Increased ms and mx, by default it was -Xms128m -Xmx256m -->
    <jvm-options>-Xms3g -Xmx3g -XX:NewSize=1g</jvm-options>
    <jvm-options>-server -Xsqnopause</jvm-options>
    <jvm-options>-XX:+UseParallelGC -XX:+UseParallelOldGC -XX:
ParallelGCThreads=20</jvm-options>

    <!-- Added java keystore locations -->
```

```

    <jvm-options>-Dcom.sun.management.jmxremote -Djavax.
net.ssl.keyStore=...keystore -Djavax.net.ssl.keyStorePassword=...</jvm-
options>
  </jvm>

  <!-- MySQL data source configured as a JDBC resource -->
  <jdbc-resource>
    <jndi-name>jdbc/rollerdb</jndi-name>
    <datasource-
class>com.mysql.jdbc.jdbc2.optional.MysqlDataSource</datasource-class>
    <property>
      <name>Password</name>
      <value>...</value>
      <description/>
    </property>
    <property>
      <name>User</name>
      <value>...</value>
      <description/>
    </property>
    <property>
      <name>URL</name>
      <value>jdbc:mysql://somedburl:12345/database?relaxAutoCommit=true&am
p;useUnicode=true&amp;characterEncoding=utf-8</value>
      <description/>
    </property>
    <description>roller datasource</description>
    <min-connections>15</min-connections>
    <max-connections>250</max-connections>
    <idle-timeout>180</idle-timeout>
    <connection-validation>meta-data</connection-validation>
  </jdbc-resource>

  <mail-resource>
    <jndi-name>mail/Session</jndi-name>
    <description>roller mail session</description>
    <property>
      <name>mail.smtp.host</name>
      <value>mailhost</value>
      <description>Mail server hostname</description>
    </property>
  </mail-resource>

  <!-- blogs web app with document root '/' runs the virtual server -->
  <virtual-server>
    <name>blogs</name>
    <host>blogs</host>
    <http-listener-name>http-listener-1</http-listener-name>
    <web-app>
      <uri>/</uri>
      <path>../web-app/blogs/_default</path>

```

```
        <description/>
    </web-app>
</virtual-server>

<!-- Logs are rotated once a day. -->
<event>
  <time>
    <time-of-day>00:00</time-of-day>
  </time>
  <rotate-access-log>true</rotate-access-log>
</event>

<event>
  <time>
    <time-of-day>00:00</time-of-day>
  </time>
  <rotate-log>true</rotate-log>
</event>
</server>
```

Sun Blogs Site Administrator's Perspective

Here's some information from the site administrator:

"As far as my experience with Sun Java System Web Server 7, I love it. We used Tomcat for the first couple years of the site when only Sun Java System Web Server 6.1 was available, and quite frankly, Sun Java System Web Server 6.1 lacked some features and ease of use. At that time, we preferred Tomcat, even though it lacked a lot of cool features.

Shortly after Sun Java System Web Server 7.0 came out, we switched and since then I've been a big fan. In my opinion, Web Server 7 offers way more features than Tomcat does, and has been very stable and nice to work with. Also, we are using Web Server because we like some of the benefits the native Web Server 7 code gives us.

One of the interesting problems we solved using Web Server 7 was an issue caused by slow clients that were tying up Web Server threads in `j2ee-module`. Basically, if a client was slow and the thread got stuck in `j2ee-module`, it would continue to tie up limited resources, such as database connections, and at one point, this became a real problem.

To resolve the problem, we utilized the ability of Web Server to set a larger `output-buffer-size` for each thread so that the response data for a slow client could be effectively buffered by Web Server and the thread wouldn't be stuck in `j2ee-module`.

In my experience, this ability is something you don't get from other web servers. I have found that Web Server 7 provides better integration between a native web server and a Java 2 Platform, Enterprise Edition (J2EE™) container than other solutions such as Apache Tomcat.

I've also been a big fan of the Web Server `perfdump` and `stats.xml` data, which can be retrieved from the Web Server at runtime, even when the site appears to be hanging from the Web. That data has been immensely valuable in debugging Web Server issues, such as the one described previously.

Apart from that, our use of Web Server 7 has been pretty much clear sailing. We've actually been running our current installations for about a year with almost no tweaks necessary. The application just sits there and runs.”

Conclusion

Sun Blogs (blogs.sun.com) offers a secure blogging service to thousands of Sun employees worldwide. The site runs the Apache Roller Weblogger application and is powered by Sun Java System Web Server 7.0, MySQL database, and Memcached, running respectively on pairs of Sun Fire T2000 and V40Z servers. The Sun Blogs site serves 4 million hits a day on average, representing one of the busiest Web 2.0 sun.com sites.

About the Authors

Meena Vyas is a software engineer with the Sun Java System Web Server development engineering team. Murthy Chintalapati was an architect and is engineering manager for Sun's Web Tier products. Allen Gilliland is a systems specialist with the sun.com engineering team that manages Sun Blogs.

For More Information

Here are additional resources:

- Sun Forums, such as Sun Java System Web Server forum:
<http://forums.sun.com/forum.jspa?forumID=759&start=0>
- Web Tier section of Sun Developer Network (SDN):
<http://developers.sun.com/webtier/>
- Sun wikis, such as:
 - Sun Java System Web Server wiki:
<http://wikis.sun.com/display/WebServer/Sun+Java+System+Web+Server>
 - Sun Java System Web Server FAQ:

<http://wikis.sun.com/display/WebServer/FAQ>

- Sun BluePrints™ wiki:
<http://wikis.sun.com/display/BluePrints/Main>
- BigAdmin wiki: <http://wikis.sun.com/display/BigAdmin/Home>
- Documentation at <http://docs.sun.com>, such as:
 - Sun Java System Web Server 7.0 Update 3 Documentation Center: <http://docs.sun.com/app/docs/doc/820-4839/fixibo?a=view>
 - *Sun Java System Web Server 7.0 Update 3 Administrator's Guide*: <http://docs.sun.com/app/docs/doc/820-4838>
 - *Sun Java System Web Server 7.0 Update 3 Performance Tuning, Sizing, and Scaling Guide*: <http://docs.sun.com/app/docs/doc/820-4846?l=en&q=Sun+Web+Server+7.0+Update+3+Performance+Tuning+and+Sizing>
- Sun Web and Proxy Servers download page:
<http://www.sun.com/download/index.jsp?cat=Web%20%26%20Proxy%20Servers&tab=3>
- Sun training courses at <http://www.sun.com/training/>, such as:
 - Sun Java System Web Infrastructure Suite Overview (WMT-ECR-1508)
 - Sun Java System Web Server 6.1: Administration and Maintenance (ECR-2307)
- Support:
 - Sun resources:
 - Register your Sun gear:
<https://inventory.sun.com/inventory/>
 - Services: <http://www.sun.com/service/>
 - SunSolveSM Online: <http://sunsolve.sun.com/>
 - Community system administration experts:
<http://www.sun.com/bigadmin/content/communityexperts/>
- Sun Java System Web Server web site:
http://www.sun.com/software/products/web_srvr/index.xml
- Related resources on BigAdmin:
 - *Using Sun Java System Web Server With Sun Identity Manager and MySQL Software*:

http://www.sun.com/bigadmin/features/articles/idm_web_mysql.jsp

- *Sun Forums: A Sun Java System Web Server 7.0 Reference Deployment:*

http://www.sun.com/bigadmin/features/articles/sun_forums_ref.jsp

- *Installing Sun Java System Web Server 7.0 on OpenSolaris Operating System:*

http://www.sun.com/bigadmin/features/techtips/websvr_opensol.jsp

- Events of interest to users of Sun products:

- Worldwide Developer Events and Sun Tech Days:

<http://developers.sun.com/events/>

- Current Events: <http://www.sun.com/events/index.jsp>

Appendix

Here is a sample of output from perfdump:

```
Sun Java System Web Server 7.0 B12/04/2006 10:15 (SunOS DOMESTIC)
```

```
Server started Fri Sep 12 09:09:07 2008
```

```
Process ... started Fri Sep 12 09:09:07 2008
```

```
ConnectionQueue:
```

```
-----
Current/Peak/Limit Queue Length      0/96/1474
Total Connections Queued              28960884
Average Queue Length (1, 5, 15 minutes) 0.00, 0.00, 0.01
Average Queueing Delay                0.05 milliseconds
```

```
ListenSocket http-listener-1:
```

```
-----
Address                http://0.0.0.0:8080
Acceptor Threads       1
Default Virtual Server blogs
```

```
KeepAliveInfo:
```

```
-----
KeepAliveCount         165/200
KeepAliveHits          13743024
```

```

KeepAliveFlushes      0
KeepAliveRefusals    1740047
KeepAliveTimeouts    2509109
KeepAliveTimeout     30 seconds

```

SessionCreationInfo:

```

-----
Active Sessions      55
Keep-Alive Sessions  0
Total Sessions Created 250/250

```

CacheInfo:

```

-----
enabled              yes
CacheEntries         275/1024
Hit Ratio            8662111/25918743 ( 33.42%)
Maximum Age         30

```

Native pools:

```

-----
NativePool:
Idle/Peak/Limit      1/1/128
Work Queue Length/Peak/Limit 0/0/0

```

DNSCacheInfo:

```

-----
enabled              yes
CacheEntries         0/1024
HitRatio             0/0 ( 0.00%)

```

Async DNS disabled

Performance Counters:

```

-----

```

	Average	Total	Percent
Total number of requests:		26058151	
Request processing time:	0.7873	20515742.0000	
default-bucket (Default bucket)			
Number of Requests:		26058151	(100.00%)
Number of Invocations:		298097638	(100.00%)
Latency:	0.0021	55890.8477	(0.27%)
Function Processing Time:	0.7852	20459852.0000	(99.73%)
Total Response Time:	0.7873	20515742.0000	(100.00%)

Sessions:

```

-----
Process Status Client      Age      VS Method  URI              Function

28793 response xxx.xx.xxx.xxx 392 blogs GET
/mramcha/resource/P1010825s.JPG
28793 response xxx.xx.xxx.xxx 58  blogs GET
/barton808/resource/MarksWelcome.JPG
28793 response xxx.xx.xxx.xxx 31  blogs GET
/katakai/entry/php_framework_support_in_netbeans
28793 response xxx.xx.xxx.xxx 28  blogs GET /wizidm/feed/entries/atom
28793 response xxx.xxx.xx.xxx 6   blogs GET /roger/ service-j2ee
28793 request  xxx.xxx.xx.xxx 4
28793 response xxx.xxx.xx.xxx 4   blogs GET /jonathan/
28793 response xx.xxx.xxx.xxx 2   blogs GET /jimgris/feed/entries/atom
28793 response xx.xxx.xxx.xxx 1   blogs GET
/PotstickerGuru/feed/entries/rss service-j2ee

...

```

Licensing Information and Third-Party Trademarks

Unless otherwise specified, the use of this software is authorized pursuant to the terms of the license found at http://www.sun.com/bigadmin/common/berkeley_license.html.

Parts of the product may be derived from Berkeley BSD systems, licensed from the University of California. UNIX is a registered trademark in the U.S. and in other countries, exclusively licensed through X/Open Company, Ltd. X/Open is a registered trademark of X/Open Company, Ltd.

All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the United States and other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.