



NOTES & DOMINO
ADMINCAMP


Extracting important Domino statistics to keep servers healthy

Andy Pedisich
Technotics

Why Do This Session?

- **Understanding statistics your servers deliver can help you to be a better administrator**
 - ♦ You can pro-actively address problems
 - ♦ You can determine answers and fixes to problems that come up by looking at historical and current data
- **The problem is, much of this data is hard to find**
 - ♦ Especially if your environment is not configured to collect it
- **This session will show you how to extract the Domino data you need to successfully run your domain**

You will find out how to make this chart in this session



What We'll Cover ...

- **Gearing up for advanced statistical analysis**
- **Creating an efficient statistic collection architecture**
- **Customizing the STATREP.NSF (Monitoring Results DB)**
- **Making sure clusters are ready for emergency failover**
- **Mastering the basics of statistical data extraction**
- **Scooping out hidden data to analyze and chart**
- **Wrap-up**

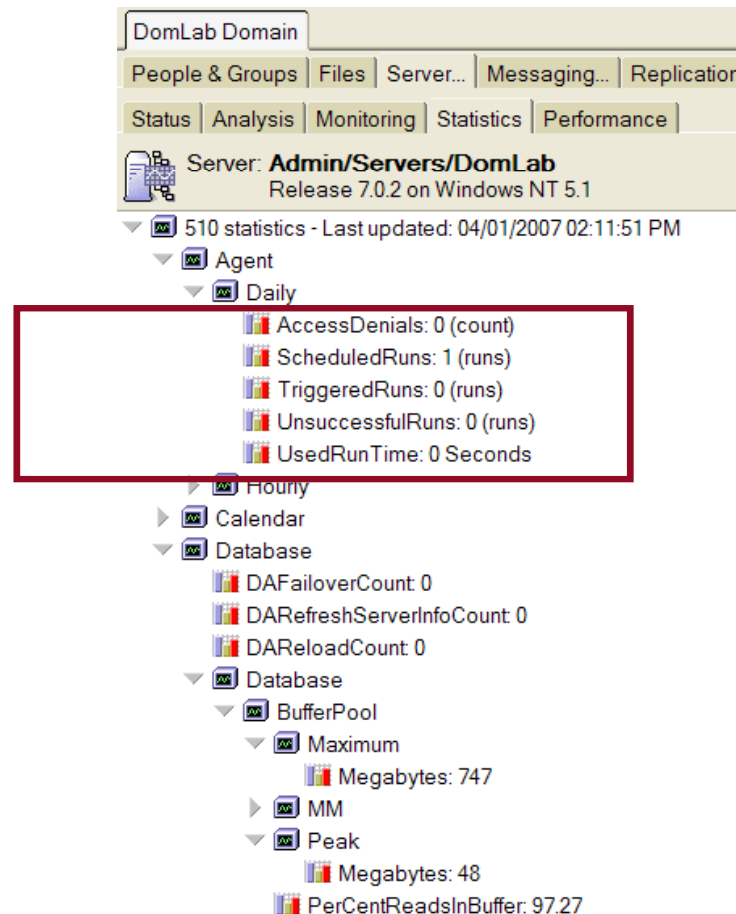
Building a Strong Foundation for Statistical Analysis

- **Statistics are generated for a broad range of categories**
 - ◆ **Domino keeps track of them hierarchically**
 - ▶ **These are the top levels of the statistical family**

ADMINP	Mem
Agent	Monitor
Calendar	NET
Database	Platform
Disk	POP3
Domino	Replica
EVENT	Server
HTTP	SMTP
LDAP	Stats
Mail	Update

Statistics Subcategory Families

- Each statistic has a multitude of sub-categories
 - ◆ This snapshot from the Administrator client shows some of the families in the statistical hierarchy
 - ▶ **“Agent” has a “Daily” sub-family, for example**



Ask for Them by Name

- Domino will produce a list of all the statistics it generates using the following console command:
 - ◆ Show stat
- You can get a list of the sub-family of a statistical hierarchy using the following console command format:
 - ◆ Show stat *top level hierarchy*
 - ◆ As in the following example:
 - ▶ **Show stat mail**



```
sh stat mail
```

```
Mail.AverageDeliverTime = 196
Mail.AverageServerHops = 1
Mail.AverageSizeDelivered = 69
Mail.CurrentByteDeliveryRate = 0
Mail.CurrentByteTransferRate = 0
Mail.CurrentMessageDeliveryRate = 0
Mail.CurrentMessageTransferRate = 0
Mail.DBCacheAged = 822
Mail.DBCacheEntries = 3
Mail.DBCacheForcedOut = 0
Mail.DBCacheHighWaterMark = 7
Mail.DBCacheHits = 8779
Mail.DBCacheMaxEntries = 504
.....
```

Drilling Down Statistics from the Console

- Use the global asterisk “*” character to see specific portions of the sub-hierarchy
 - ◆ You can use SH instead of SHOW
 - ◆ If you only want Server.Users hierarchy, use the global “*”
 - ▶ **Show stat mail.deliveredsize.***
 - *Or you can abbreviate the console command*
 - Sh stat mail.deliver*

```
sh stat mail.deliveredsize.*
```

```
Mail.DeliveredSize.100KB_to_1MB = 102  
Mail.DeliveredSize.10KB_to_100KB = 1837  
Mail.DeliveredSize.1KB_to_10KB = 7730  
Mail.DeliveredSize.1MB_to_10MB = 21  
Mail.DeliveredSize.Under_1KB = 41  
5 statistics found
```

```
sh stat mail.deliver*
```

```
Mail.Delivered = 9731  
Mail.DeliveredSize.100KB_to_1MB = 102  
Mail.DeliveredSize.10KB_to_100KB = 1837  
Mail.DeliveredSize.1KB_to_10KB = 7730  
Mail.DeliveredSize.1MB_to_10MB = 21  
Mail.DeliveredSize.Under_1KB = 41  
Mail.Deliveries = 10329  
Mail.DeliveryThreads.Active = 0  
Mail.DeliveryThreads.Max = 8  
Mail.DeliveryThreads.Total = 3  
10 statistics found
```



Secret

The Four General Types of Statistics

1. Stats that don't change after starting a Domino server

- ◆ **Disk.D.Size = 71,847,784,448**
- ◆ **Server.Version.Notes = Release 8.5.2**
 - ▶ **These have great value because they can be used to help produce an inventory of server platform configurations**

2. Stats that are snapshots of a current moment

- ◆ **Mem.Free = 942,411,776**
- ◆ **Server.AvailabilityIndex = 83**
 - ▶ **Valuable because they show what's happening now**
 - ▶ **Can be used for problem determination**

The Four General Types of Statistics (cont.)

3. Stats that are calculated

- ◆ SMTP.SessionDuration.Ave = 4
- ◆ SMTP.SessionDuration.Max = 593
 - ▶ **They are calculated from the time a server starts**
 - *These can help you understand trends*

4. Stats that are cumulative

- ◆ Mail.TotalKBTransferred.SMTP = 4,267,804,067
- ◆ Server.Trans.Total = 394385
 - ▶ **They are also accumulated from the time the server starts**
- ◆ But they only have value if extracted and analyzed
 - ▶ **Which is precisely what we are going to do**



Note

What We'll Cover ...

- Gearing up for advanced statistical analysis
- **Creating an efficient statistic collection architecture**
- Customizing the STATREP.NSF (Monitoring Results DB)
- Making sure clusters are ready for emergency failover
- Mastering the basics of statistical data extraction
- Scooping out hidden data to analyze and chart
- Wrap-up

The Two Things Needed

- **Two things are required for statistics collection:**
 - ◆ **The Collect task must be running on any server that is designated to collect the statistics**
 - ▶ **Not all servers should run the Collect task**
 - ◆ **The EVENTS4 database must have at least one Statistics Collection document**
 - ▶ **Statistics should be collected centrally on one or two servers so that the data is easy to get to**
 - *Stats should be collected every hour to be effective*
 - ▶ **EVENTS4 should be the same replica on all servers in the domain**



Checklist

We Know What the Replica ID Should Be for EVENTS4

- The replica ID of system databases, such as EVENTS4, is derived from the replica ID of the address book

<u>Database</u>	<u>Replica ID</u>
NAMES.NSF	852564AC:004EBCCF
CATALOG.NSF	852564AC:014EBCCF
EVENTS4.NSF	852564AC:024EBCCF
ADMIN4.NSF	852564AC:034EBCCF

- ▶ Notice that the first two numbers after the colon for the EVENTS4.NSF replica are 02



Don't
Forget

- *Make sure that EVENTS4.NSF is the same replica ID throughout the domain by opening it and putting it on your desktop*

Want to Add Every EVENTS4.NSF to Your Desktop?

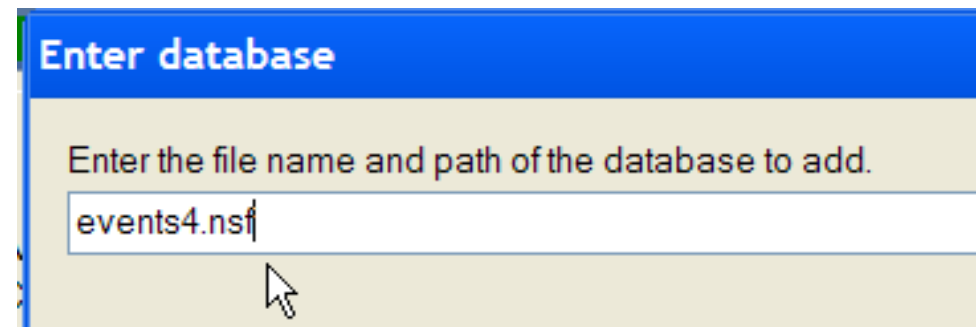
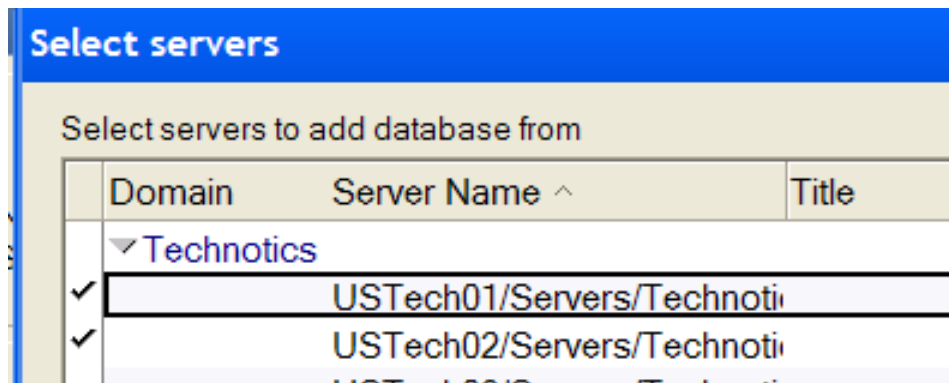
- Add this code to a button on your toolbar
 - ◆ This is courtesy of Thomas Bahn
 - ▶ www.assono.de/blog



```
_names := @Subset(@MailDbName; 1) : "names.nsf";  
  
_servers := @PickList([Custom]; _names; "Servers"; "Select  
servers"; "Select servers to add database from"; 3);  
  
_db := @Prompt([OkCancelEdit]; "Enter database"; "Enter the file  
name and path of the database to add."; "log.nsf");  
  
@For( n := 1; n <= @Elements(_servers); n := n + 1;  
@Command([AddDatabase]; _servers[n] : _db) )
```

Add a Database to the Desktop

- This code will prompt you to pick the servers that have the database you want on your desktop
 - ◆ Then it will prompt for the name of the database
 - ▶ **And open it on all the servers you've selected**
- Use it to make sure all the EVENTS4.NSF are the same replica in your domain



A Required Design, but No Required Name

- There has to be a Statrep.nsf on every server
 - ♦ It is used by the server to store monitoring data
 - ▶ **It must be designed using the Statrep5.ntf Monitoring Results template**
 - *Its default title is Monitoring Results*
- But you don't have to use one of those for your statistic collection repository
 - ♦ Create your own collection points and give the database a unique name

City	Collecting Server	Monitoring Results DB
New York	USNYAdmin1	USStatrep.nsf
Amsterdam	EUNeHub01	EUStatrep.nsf

What We'll Cover ...

- Gearing up for advanced statistical analysis
- Creating an efficient statistic collection architecture
- Customizing the STATREP.NSF (Monitoring Results DB)
- Making sure clusters are ready for emergency failover
- Mastering the basics of statistical data extraction
- Scooping out hidden data to analyze and chart
- Wrap-up

Why Would the Statrep.nsf Need Customization?

- The Statrep database is surprisingly lacking in value
- There are eight views of statistical data in Statrep
 - ◆ The Calendaring & Scheduling and Clusters views lack any substantial information
 - ◆ The Communications view strictly shows dial-up stats
 - ▶ **What year is this?**

Collection Time	Server Name	Total Users	Total Appointment	Total Reservations
▼ Admin/Servers/DomLab				
04/14/2007 02:12:08 PM	Admin/Servers/DomLab	7	0	1
04/13/2007 01:18:17 PM	Admin/Servers/DomLab	7	0	1

Collection Time	Server Name	Cluster Name
▼ USAppDev01/TheViewLab		
06/28/2004 10:08:16 AM	USAppDev01/TheViewLab	AppCluster1
06/28/2004 09:07:51 AM	USAppDev01/TheViewLab	AppCluster1

Collection Time	COM1 Errors: CRC	Port	Retransmitted Packets
▼ USAppDev01/TheViewLab			
06/28/2004 10:08:16 AM	N/A	N/A	N/A
06/28/2004 09:07:51 AM	N/A	N/A	N/A



More Marginal Views

- The Mail & Database view is not bad
- The Network view is useless
- The Platform and System views have some OK data
 - ◆ But they still contain some questionable data

Collection Time	Dead Mail	Mail Routed	Pend. Mail	Failed Replications	Successful Replications
USTech01/Servers/Technotics					
04/07/2007 04:29:53 PM	228	11,313	1	72	288
04/07/2007 03:29:53 PM	228	11,298	1	71	284

Collection Time	Port 1	Port 2	Port 3	Port 4	Port 5	Port 6
USTech01/Servers/Technotics						
04/07/2007 04:29:53 PM	TCPIP	NA	NA	NA	NA	NA
04/07/2007 03:29:53 PM	TCPIP	NA	NA	NA	NA	NA
04/07/2007 02:29:53 PM	TCPIP	NA	NA	NA	NA	NA

Collection Time	Platform	Total Disk Util	Memory Free	Memory Pages/Sec	Memory Scan Rate	Memory Util	Total Network Bytes	Total Network Util	Total Process	Total CPU Util
USAppDev01/TheViewLab										
06/28/2004	Windows NT	N/A	105.5	0	N/A	79	42782.9	0	3.4	7.4
06/28/2004	Windows NT	N/A	113.9	0	N/A	77	20354.1	0	2.8	5.5

Collection Time	Space on Data Path	Swap File/Sysvol Size	Dead Mail	Pend. Mail	Users	Mem Alloc'd	Trans/Min
USTech01/Servers/Technotics							
04/07/2007	43,495,075,840	NA	228	1	6	323,942,812	51
04/07/2007	43,495,264,256	NA	228	1	4	323,695,902	68

Web Server Stats Documents

- **Web server stats documents are strangely empty**
 - ◆ **Mostly because they reference stats that aren't there**
 - ▶ **And I need views to really get trend information!**

Web Server Statistics	
Web server version:	
Total number of login commands received:	
HTTP requests in last minute:	
HTTP requests in last 5 minutes:	
HTTP requests in last hour:	
Total number of commands received:	2,854
Total number of unknown commands received:	



~~Domino.Command.GetOrbCookie = 0~~
 Domino.Command.MoveToFolder = 0
 Domino.Command.Navigate = 0

 Domino.Command.OpenWebService = 0
 Domino.Command.ReadForm = 74
 Domino.Command.Redirect = 0
 Domino.Command.RemoveFromFolder = 0
 Domino.Command.RequestCert = 0
 Domino.Command.SaveDocument = 0

Web Server Statistics	
Web server version:	Domino.BuildVersion T
Total number of login commands received:	Domino.Command.Login #
HTTP requests in last minute:	Domino.Requests.Per1Minute #
HTTP requests in last 5 minutes:	Domino.Requests.Per5Minute #
HTTP requests in last hour:	Domino.Requests.Per1Hour #
Total number of commands received:	Domino.Command.Total #
Total number of unknown commands received:	Domino.Command.Uknown #

~~Domino.Command.RequestCert = 0~~
~~Domino.Command.SaveDocument = 0~~

 Domino.Requests.Per1Day.Peak = 854
 Domino.Requests.Per1Day.PeakTime = 03/08/2006 09:39:24 MST
~~Domino.Requests.Per1Day.Total = 77~~
 Domino.Requests.Per1Hour.Peak = 339
 Domino.Requests.Per1Hour.PeakTime = 03/07/2006 18:29:39 MST
 Domino.Requests.Per1Hour.Total = 0
 Domino.Requests.Per1Minute.Peak = 70
 Domino.Requests.Per1Minute.PeakTime = 03/07/2006 18:27:29 MST
 Domino.Requests.Per1Minute.Total = 0
 Domino.Requests.Per5Minute.Peak = 162
 Domino.Requests.Per5Minute.PeakTime = 03/07/2006 18:27:29 MST

Web Server Stats

- If you are running Domino Web Access, you're probably very interested in how your Web servers are doing
 - ◆ You'd like to know how much work is being done
- But the Web Server/Web Retriever view is also useless
 - ◆ The columns reference fields that don't exist!

	Collection Time	Server Name	Web Server Logins	Requests in Last Hour	Web Retriever Requests
	▼ USTech01/Servers/Technotics				
★	03/12/2006 03:08:00 PM	USTech01/Servers/Technotics			
★	03/12/2006 02:08:00 PM	USTech01/Servers/Technotics			
★	03/12/2006 01:08:00 PM	USTech01/Servers/Technotics			



Warning

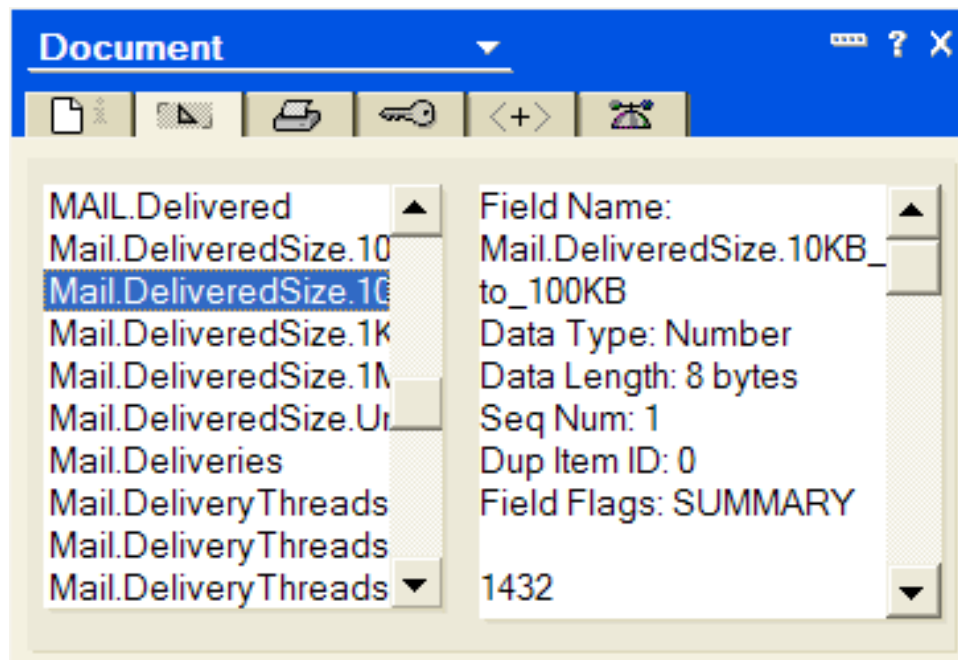
Documents Have Better Content Than Views

- In most cases, the documents in the Statrep database have much better content than the views
 - ◆ The statistics in the document are just really hard to use
 - ▶ You really need the context of seeing many samples in a column format to make sense out of it

Notes Mail Statistics	
Number of dead mail messages:	228
Total number of mail messages routed:	11,223
Mail waiting to be routed:	3
Total number of mail messages delivered:	7,520
Mail waiting to be delivered:	8
Minimum delivery time per message:	1 seconds
Average delivery time per message:	20 seconds
Maximum delivery time per message:	12,891 seconds
Minimum size of message delivered:	1 KB
Average size of message delivered:	71 KB
Maximum size of message delivered:	6,146 KB

Everything Is Everywhere

- Keep in mind that every single statistic that is generated is contained in every document in the Monitoring Results database
 - ◆ You just can't see it all because it's not in views or documents
 - ▶ **And views are the most important place to have it because that's where it gives you the ability to compare samples**
 - *And analyze trends*



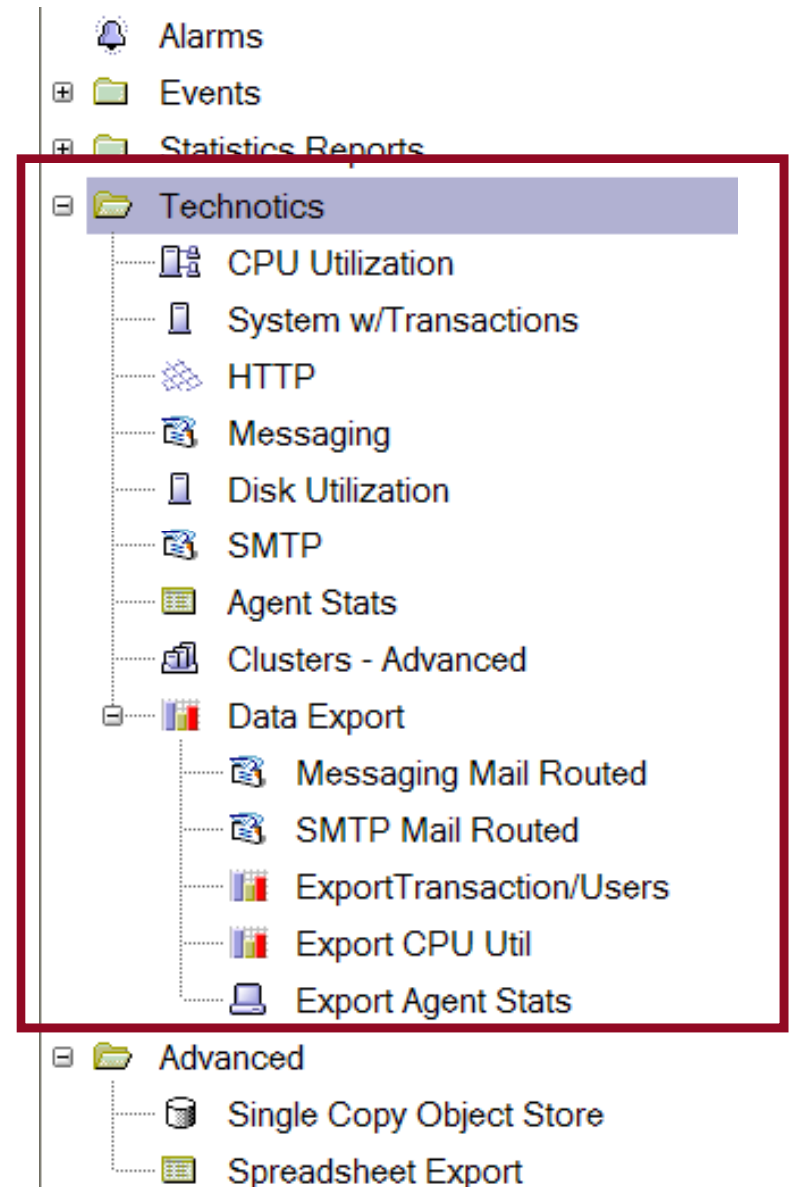
Secret

The Stats Are There, Now You See Them

- There is a new, customized version of the Monitoring Results database on your conference CD called:
 - ◆ TechnoticsStatrep10-9.NSF
- It has all the views that are on the original Statrep
 - ◆ Plus over a dozen additional views to help you analyze statistics from your servers

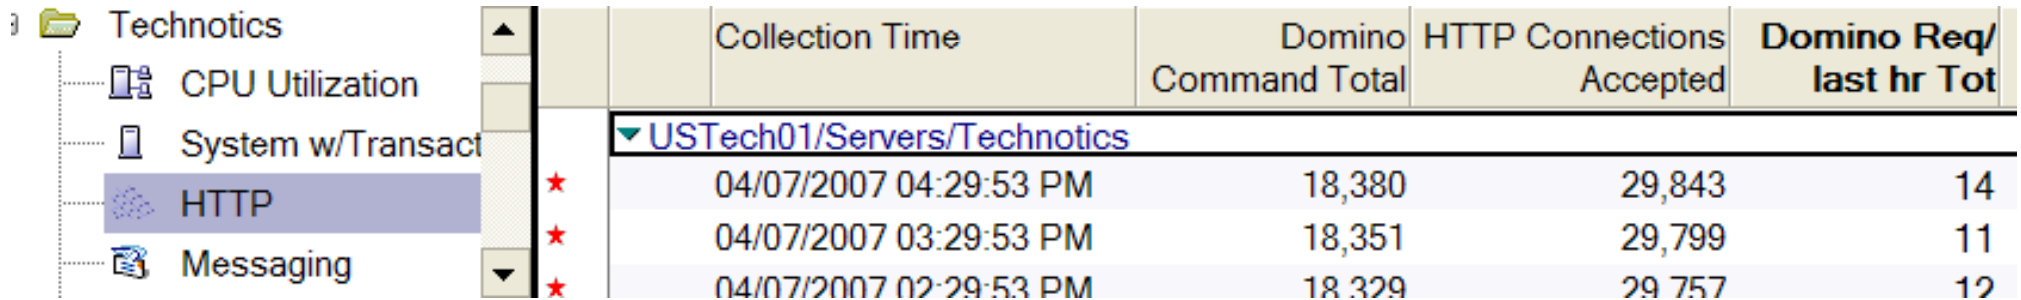


Take Home



Here's a View You Can Use

- That same TechnoticsStatrep10-9.NSF has a handy view that's been modified to help you see what's happening on the web



The screenshot shows a Domino database view for the Technotics folder. The left sidebar contains a tree view with the following items: Technotics (selected), CPU Utilization, System w/Transact, HTTP (highlighted), and Messaging. The main table displays data for the selected folder, with columns for Collection Time, Domino Command Total, HTTP Connections Accepted, and Domino Req/last hr Tot. The data is organized into a hierarchy under USTech01/Servers/Technotics, showing three rows of data for different collection times on 04/07/2007.

Collection Time	Domino Command Total	HTTP Connections Accepted	Domino Req/last hr Tot
▼ USTech01/Servers/Technotics			
04/07/2007 04:29:53 PM	18,380	29,843	14
04/07/2007 03:29:53 PM	18,351	29,799	11
04/07/2007 02:29:53 PM	18,329	29,757	12

Even Platform Statistics Are Available

- The Platform Statistics view will show important aspects of the OS and platform supporting your server
 - ◆ Keep in mind that these platform statistics are momentary snapshots produced by the underlying OS
 - ▶ Their sample rates are different than the performance monitoring tools available with your platform
 - *Therefore the results will not always be exactly the same as your platform results*
 - ◆ The view adds all updall and agent manager threads

	Collection Time		Av Inx	Domino CPU	Server	RAM Util	Trans /Min	Users	All Upd	All Amgr	RAM Free	Pages /Sec	Trans total	
Alarms														
Events														
Statistics Reports		▼ USTech01												
Technotics	★	04/07/2007 04:29:53 PM	Windows NT 5.0	100	1.5	0.3	86	51	6	0.1	0.0	69	17.7	321,184
CPU Utilization	★	04/07/2007 03:29:53 PM	Windows NT 5.0	100	0.8	0.3	88	68	4	0.0	0.0	57	4.4	317,801
System w/Trans	★	04/07/2007 02:29:53 PM	Windows NT 5.0	100	0.6	0.2	74	4	2	0.0	0.0	129	14	314,325
	★	04/07/2007 01:29:53 PM	Windows NT 5.0	100	4.3	0.3	70	60	2	3.3	0.0	149	177.8	311,170

Disk Statistics Are Available, Too

- The views focus on two important aspects of disk performance
 - ◆ Disk utilization and disk queue length
 - ▶ High disk queue lengths can indicate hardware issues
 - *Not all platforms produce the same statistics for disk information*

Collection Time	D1 Util	D1 Util Avg	D1 QLen	D1 QLen Avg	D2 Util	D2 Util Avg	D2 QLen	D2 QLen Avg	D3 Util	D3 Util Avg	D3 QLen	D3 QLen Avg	D4 Util	D5 Util	D6 Util
▼ USMail05/TheViewLab															
* 06/28/2004 10:12:20 AM	1.0	2.5			0.3	2.5			0.0	0.0			0.0	0.0	4.4
* 06/28/2004 09:11:48 AM	0.7	2.4			0.2	2.4			0.0	0.0			0.0	0.0	4.3
* 06/28/2004 08:11:58 AM	1.0	2.4			0.4	2.4			0.0	0.0			0.0	0.0	2.2
* 06/28/2004 07:11:00 AM	1.0	2.4			0.4	2.4			0.0	0.0			0.0	0.0	1.5

Plus, a View to Help Analyze Shrinking Disk Space

- Showing free disk space by the hour will help you determine the cause of extreme disk space usage
 - ◆ This view is built for Domino running on Wintel

Help			
	Collection Time	Disk C: Free	Disk D: Free
▼ USMail01/TheViewLab			
*	02/08/2009 07:40:34 PM	4.5G	24.5G
*	02/08/2009 06:40:34 PM	4.5G	24.2G
*	02/08/2009 05:40:34 PM	4.5G	24.2G
*	02/08/2009 04:40:34 PM	4.5G	24.2G
*	02/08/2009 03:40:34 PM	4.5G	24.2G
*	02/08/2009 02:40:34 PM	2.9G	3.8G
*	02/08/2009 01:40:34 PM	2.9G	3.8G
*	02/08/2009 12:40:35 PM	2.9G	3.8G
*	02/08/2009 11:40:34 AM	2.9G	1000.0M
*	02/08/2009 10:40:34 AM	2.9G	1008.6M
*	02/08/2009 09:40:34 AM	2.9G	1.0G
*	02/08/2009 08:40:34 AM	2.9G	1.1G
*	02/08/2009 07:40:34 AM	2.9G	1.1G
*	02/08/2009 06:40:34 AM	2.9G	1.1G
*	02/08/2009 05:40:34 AM	2.9G	1.1G
*	02/08/2009 04:40:34 AM	2.9G	1.1G

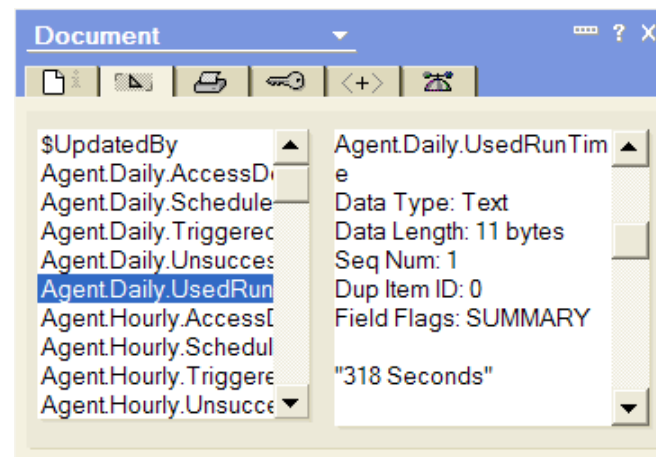
Demo: Checking Out the Views of Technotics STATREP

Demo



Caveats About the Technotics Statrep

- Some views expose the following statistics
 - ◆ Agent.Daily.UsedRunTime and Agent.Hourly.UsedRunTime
 - ▶ This stat generated the agent runs in seconds
- Some versions of Domino produce this stat as a text field, others as a numeric field
 - ◆ A formula converts it to a numeric field
 - ▶ This might not be necessary in your domain
 - `@If(@IsAvailable(Agent.Hourly.UsedRunTime);(@TextToNumber(@LeftBack(Agent.Hourly.UsedRunTime;8))/60);"N/A")`
 - ▶ The formula also converts the statistic from seconds to minutes



One More Caveat

- **A few views display disk utilization statistics such as:**
 - ◆ **Platform.LogicalDisk.2.AvgQueueLen.Avg**
- **Disk statistic names vary from platform to platform**
 - ◆ **AIX and iSeries systems can have much longer device names**
 - ◆ **Even in the Wintel platform they can be listed as:**
 - ▶ **Logical disks**
 - ▶ **Physical disks**
- **Be sure to check Statrep to see how it is represented in your domain**
 - ◆ **You might find it necessary to customize all disk views for your own environment**

What We'll Cover ...

- **Gearing up for advanced statistical analysis**
- **Creating an efficient statistic collection architecture**
- **Customizing the STATREP.NSF (Monitoring Results DB)**
- **Making sure clusters are ready for emergency failover**
- **Mastering the basics of statistical data extraction**
- **Scooping out hidden data to analyze and chart**
- **Wrap-up**

Cluster Replication Basics

- Cluster replication keeps the database on the primary server in sync with the replica on the failover server
 - ♦ Cluster replication is an event-driven process that occurs automatically when a change is made to a database
 - ▶ **It's vital that these replicas are synchronized**
 - *But by default, servers in a cluster only have a single cluster replicator thread between them*



Problem

Can the Single Cluster Replicator Keep Up?

- Occasionally there is too much data changing to be replicated efficiently by a single cluster replicator
 - ◆ If cluster replicators are too busy, replication is queued until more resources are available and databases get out of sync
 - ▶ Then a database on a failover server does not have all the data it's supposed to have
- If users must failover to a replica on a different server, they think their information is gone forever!
 - ◆ All because replicas will not have the same content
 - ▶ Users need their cluster insurance!



Caution

How Many Is Enough?

- Adding a cluster replicator will help fix this problem
 - ◆ Use this parameter in the Notes.ini
 - ▶ **CLUSTER_REPLICATORS=#**
 - ◆ Add one dynamically from the console using this command
 - ▶ **Load clrepl**
- The challenge is to have enough cluster replicators without adding too many
 - ◆ Adding too many clusters will have a negative effect on server performance
- Here are some important statistics to watch so that you can make a wise decision about how many to add!



Decision
Point

Key Stats for Vital Information About Cluster Replication

Statistic	What It Tells You	Acceptable Values
Replica.Cluster.SecondsOnQueue	Total seconds that last DB replicated spent on work queue	< 15 sec – light load < 30 sec – heavy
Replica.Cluster.SecondsOnQueue.Avg	Average seconds a DB spent on work queue	Use for trending
Replica.Cluster.SecondsOnQueue.Max	Maximum seconds a DB spent on work queue	Use for trending
Replica.Cluster.WorkQueueDepth	Current number of databases awaiting cluster replication	Usually zero
Replica.Cluster.WorkQueueDepth.Avg	Average work queue depth since the server started	Use for trending
Replica.Cluster.WorkQueueDepth.Max	Maximum work queue depth since the server started	Use for trending

What to Do About Stats Over the Limit

- **Acceptable Replica.Cluster.SecondsOnQueue**
 - ◆ Queue is checked every 15 seconds, so under light load should be less than 15
 - ▶ Under heavy load, if the number is larger than 30, another cluster replicator should be added
- If the above statistic is low and **Replica.Cluster. WorkQueueDepth** is constantly higher than 10 ...
 - ◆ Perhaps your network bandwidth is too low
 - ▶ Consider setting up a private LAN for cluster replication traffic



**Best
Practice**

The Documents Have More Information

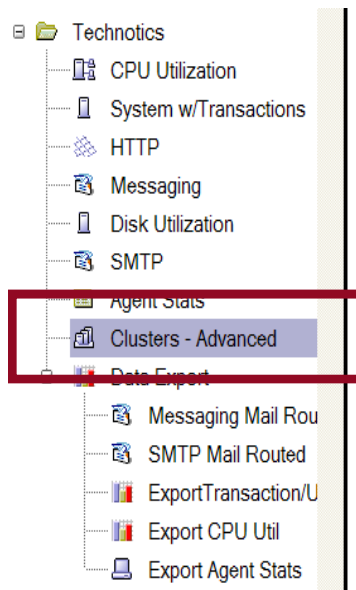
- The cluster documents have much better information than the default cluster views
 - ◆ But they still lack key stats, although they are in each doc

Server Cluster Statistics	
Default port name:	TCPIP
Server availability index:	17
Server availability threshold:	10
Successful redirects:	0
Unsuccessful redirects:	0
Successful redirects by path:	0
Unsuccessful redirects by path:	0
Successful redirects if server busy:	0
Unsuccessful redirects if server busy:	0
Successful redirects by path if server busy:	53
Unsuccessful redirects by path if server busy:	0
Database open requests where all servers in cluster are busy:	0
Database open requests where database is marked out of service:	0
Database open requests where this server is busy:	9580
Number of probes sent to other servers in cluster:	2454
Number of probe errors:	0

Replica Cluster Statistics	
Number of servers in replica cluster	1
Number of successful cluster replications	4633
Number of unsuccessful cluster replications:	
Number of docs added:	1092
Number of docs updated:	2774
Number of docs deleted:	480
Number of cluster replicas on this server:	795
Number of cluster replicas on other servers in this cluster:	796

Statistics Missing from the STATREP that comes with Notes

- The Technotics Statrep tracks the key statistics you need to help track and adjust your clusters
 - ◆ It also has a column for the Server Availability Index



Take Home

Collection Time	Server Name	Cluster Na	Exp Fact	Av Inx	Min on Q	Min/Q Av	Min/Q Mx	WkrDpth	WD Av	WD Mx
▼ Demo4/Servers/TheView										
07/20/2004 11:40:46 AM	Demo4/Servers/TheView	D4Cluster	26.0	22	10.3	2.8	210.8	165	42	1294
07/20/2004 11:40:43 AM	Demo4/Servers/TheView	D4Cluster	26.0	22	10.3	2.8	210.8	165	42	1294
07/20/2004 10:40:22 AM	Demo4/Servers/TheView	D4Cluster	15.1	35	1.1	2.8	210.8	21	42	1294
07/20/2004 10:40:21 AM	Demo4/Servers/TheView	D4Cluster	15.1	35	1.1	2.8	210.8	21	42	1294
07/20/2004 09:39:58 AM	Demo4/Servers/TheView	D4Cluster	17.8	31	31.9	2.7	210.8	284	42	1294
07/20/2004 09:39:56 AM	Demo4/Servers/TheView	D4Cluster	17.8	31	31.9	2.7	210.8	284	42	1294

My Column Additions to Statrep

Column Title	Formula	Formatting
Min on Q	Replica.Cluster.SecondsOnQueue/60	Fixed (One Decimal Place)
Min/Q Av	Replica.Cluster.SecondsOnQueue.Avg/60	Fixed (One Decimal Place)
Min/Q Mx	Replica.Cluster.SecondsOnQueue.Max/60	Fixed (One Decimal Place)
WkrDpth	Replica.Cluster.WorkQueueDepth	General
WD Av	Replica.Cluster.WorkQueueDepth.Avg	General
WD Mx	Replica.Cluster.WorkQueueDepth.Max	General

Demonstration: Looking at Cluster Statistics



Demo

What We'll Cover ...

- **Gearing up for advanced statistical analysis**
- **Creating an efficient statistic collection architecture**
- **Customizing the STATREP.NSF (Monitoring Results DB)**
- **Making sure clusters are ready for emergency failover**
- **Mastering the basics of statistical data extraction**
- **Scooping out hidden data to analyze and chart**
- **Wrap-up**

The Statrep Template's Only Export View

- The default Lotus Statrep template's Spreadsheet Export view just doesn't seem to give us enough power
 - ◆ Pulling the data into Microsoft Excel, then analyzing and graphing the data can often give you amazing insight into usage patterns
 - ▶ **This information will be invaluable when:**
 - *Trying to consolidate servers*
 - *Troubleshooting performance issues*

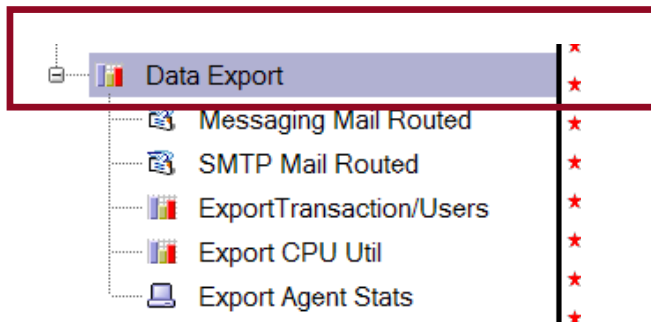


Tip

Monitoring Results		Help						
			Space on Data Path	Swap File/ Sysvol Si	Users	Trans/Min	Mem Free	
⊕ Statistics Reports	▲	01/Servers/Tec1	04/06/2007 08:59:53 AM	43,599,925,248	6	1	1,009,119,232	
⊕ Technotics		01/Servers/Tec1	04/06/2007 07:59:52 AM	43,601,068,032	5	0	1,017,745,408	
⊖ Advanced		01/Servers/Tec1	04/06/2007 06:59:52 AM	43,600,756,736	4	0	1,020,264,448	
Single Copy Object Store		01/Servers/Tec1	04/06/2007 05:59:52 AM	43,595,493,376	4	0	1,033,990,144	
Spreadsheet Export	▼	01/Servers/Tec1	04/06/2007 04:59:52 AM	43,595,558,912	4	0	1,036,865,536	
		01/Servers/Tec1	04/06/2007 03:59:51 AM	43,595,935,744	4	0	1,030,955,008	

Analysis Tools

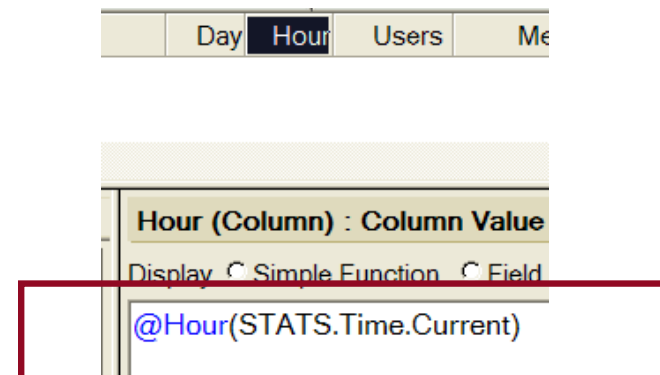
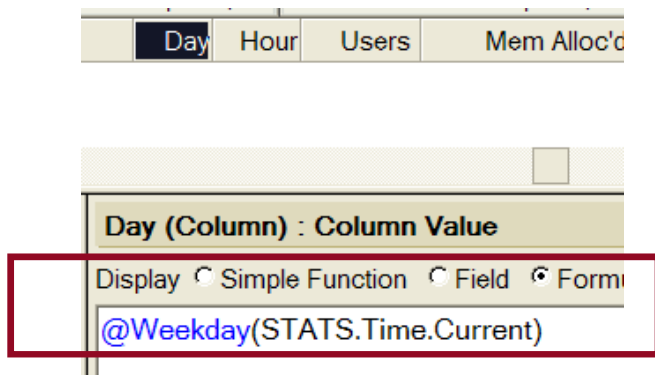
- Let's cover the basics of the Statrep views used in the data export process
 - ◆ And a special Excel spreadsheet that contains custom formulas



Master Export Formulas					
CPU Util Stats Spreadsheet					
Trans total /Hr					
Copy this cell---->	0				
Transactions and Users Spreadsheet					
Trans/Hour					
Copy this cell---->	#VALUE!				
SMTP Messages Spreadheet					
	SMTP Mess Processed/Hr	SMTP Size Tot KB/Hr	Mess Recip Tot/Hr	Mail Tot KB Trans/hr	Mail Tot Routed SMTP/Hr
Copy these cells---->	0	0	0	0	#VALUE!

You Need a Better View of the Situation

- The data export views are designed to be exported as CSV files
 - ◆ Each has key fields that are important to the export
 - ▶ **Hour and Day generate an integer that represents the hour of the day and a day of the week**
 - *Hour 15 = 3:00 PM*
 - *Day 1 = Sunday, Day 7 = Saturday*
 - *These are used in hourly and daily calculations in pivot tables*



Export Views Are All Flat Views

- Any view that is used for exporting data is flat, not categorized
 - ◆ This makes it easier to manipulate in pivot tables in Excel
- There are columns in the export views that appear to have no data
 - ◆ They will be filled with a formula when brought into Excel

Server	Collection Time	Trans total	Minutes of Agent Run Last Hour	Trans total /Hr
USMail03	06/15/2004 06:08:41 PM	17,288	0	
USMail03	06/15/2004 05:08:53 PM	19,160	0	
USMail03	06/15/2004 01:30:29 PM	22,419	0	
USMail03	06/15/2004 12:30:08 PM	23,909	0	
USMail03	06/15/2004 11:29:57 AM	24,684	0	
		25,453	0	

Formulas Are Already Available

- There is a spreadsheet containing my formulas to help you develop charts for all of this data
 - ◆ Master Formula XLS Domino Stat Exports- Technotics -V 2-4.xls
 - ▶ One place to find this spreadsheet is on my blog
 - ◆ <http://www.andypedisich.com/blogs/andysblog.nsf/dx/resources.htm>
- The views and this spreadsheet will all fit together in a few moments



What We'll Cover ...

- **Gearing up for advanced statistical analysis**
- **Creating an efficient statistic collection architecture**
- **Customizing the STATREP.NSF (Monitoring Results DB)**
- **Making sure clusters are ready for emergency failover**
- **Mastering the basics of statistical data extraction**
- **Scooping out hidden data to analyze and chart**
- **Wrap-up**

Transactions per Hour

- This can be a very important statistic if you are thinking about consolidation
 - ◆ Use time span to sample all servers for the best results
 - ▶ **It will allow you to compare apples to apples**
 - ◆ And because all the export data contains a reference to the day of the week, you could select the data for Monday through Friday to get the significant averages



**Don't
Forget**

Examining Transactions

- If a few servers are performing badly, you might want to know how many transactions they are processing
 - ◆ Especially if the servers have the same hardware
 - ▶ **And if they have a similar number of mail users assigned**
- I want to compare these servers statistically
 - ◆ What I want to know is:
 - ▶ **How many users are hitting these systems?**
 - ▶ **How many transactions are these servers being forced to make?**
 - ***And I want to know these things on a PER HOUR basis***



Start by Going to the Export Transactions/Users View

- Analysis starts with Export Transactions/Users view
 - ◆ I don't hesitate to add new views to Statrep
 - ▶ I don't change the old ones, I just add new ones
- Note that Trans/Total is a cumulative stat
 - ◆ And the Trans/Hour column is blank
 - ▶ We have a custom formula to apply to this column after the data is exported into MS Excel

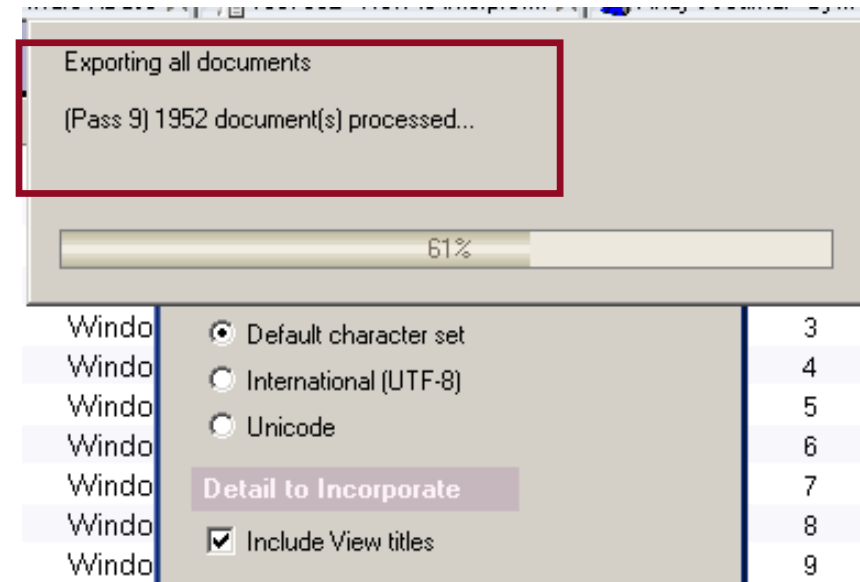
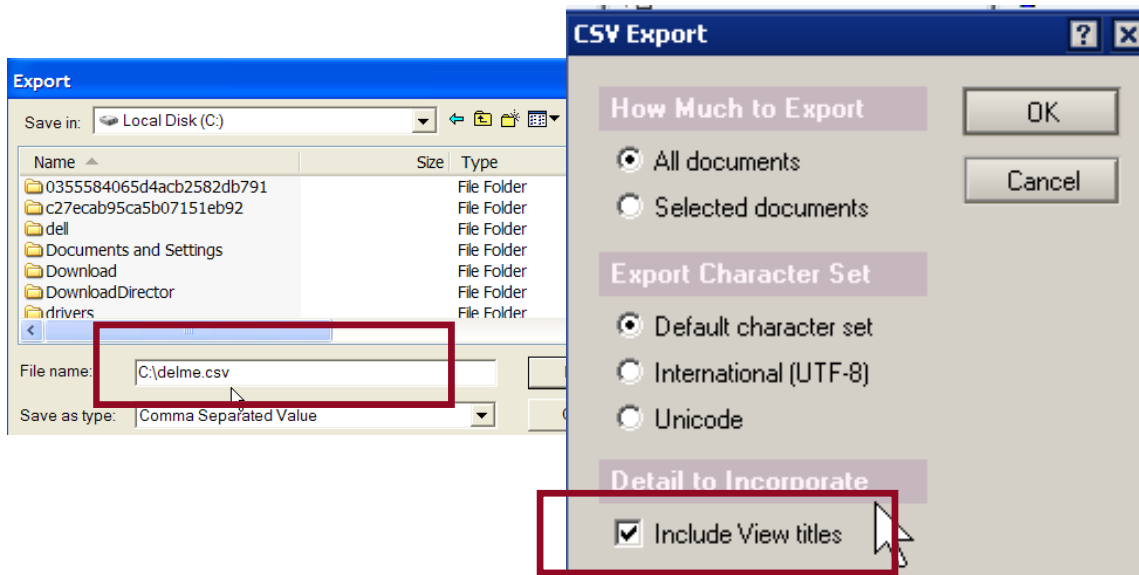
Server	OS	Collection Time	Day	Hour	Users	Mem Alloc'd	Trans/Min	Trans/Total	Trans/Hour
USAppDev01	Windows NT 5.0	06/27/2004 07:08:01 AM	1	7	3	222,646,026	188	302,578	



Secret

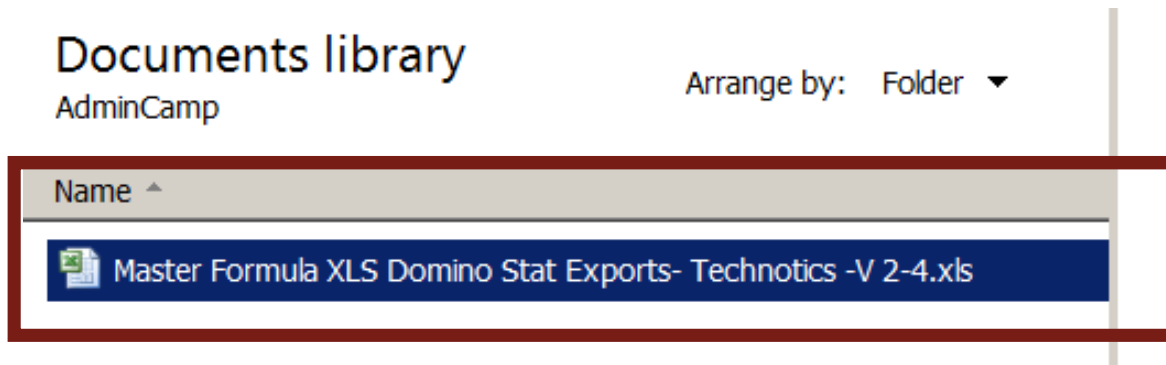
Next, Export View to CSV File

- I export the contents of the view to a CSV file
 - ◆ Before Vista/Windows 7, the file is always called C:\delme.csv
 - ▶ Vista and Windows 7 won't let you put a file in C:\
 - ▶ You can now put it with your personal documents
 - ◆ Don't forget to include the view titles
 - The import is very fast, even when there is a lot of data



Next, Open the Special Spreadsheet

- Start Excel and open the spreadsheet containing the formulas to help you develop charts for all of this data
 - ◆ Master Formula XLS Domino Stat Exports- Technotics -V 2-4.xls



What's in the Spreadsheet?

- The spreadsheet contains the formulas that will help to break down server activity into per hour averages
 - ◆ Don't worry about the #value errors
 - ▶ Then open the DELME.CSV file

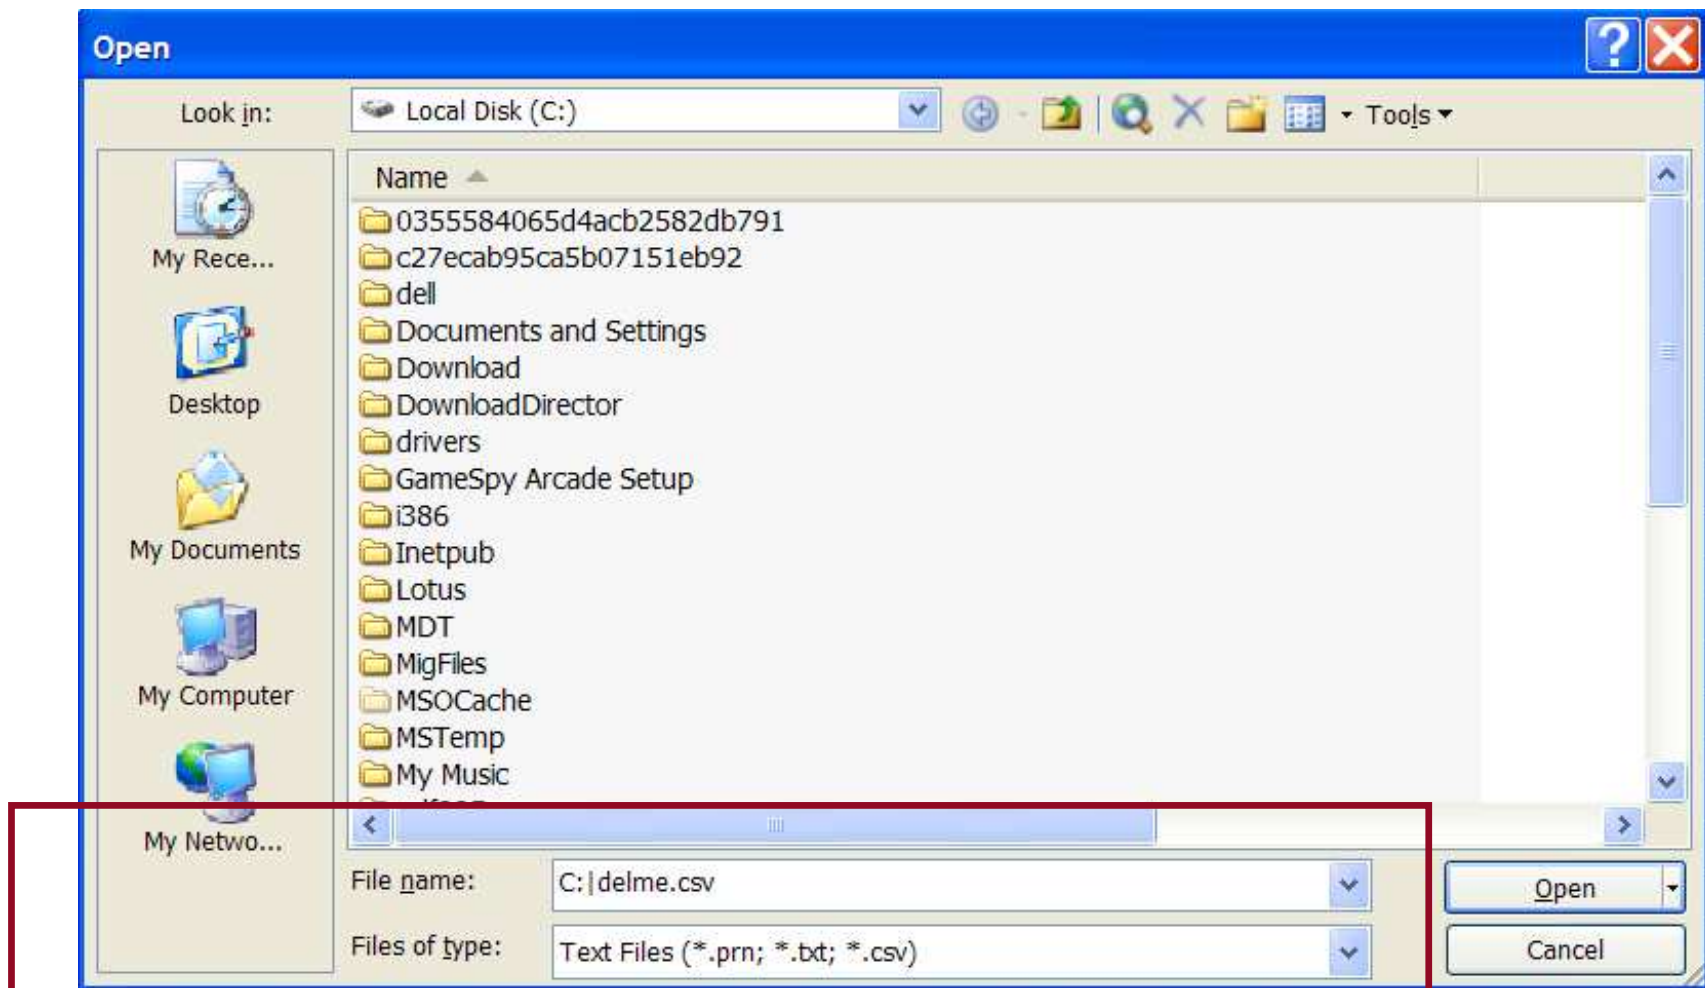
S	T	U	V	W	X
Technotics, Inc.					
Master Export Formulas					
CPU Util Stats Spreadsheet					
	Trans total /Hr				
Copy this cell---->	0				
Tranasactions and Users Spreadsheet					
	Trans/Hour				
Copy this cell---->	#VALUE!				
SMTP Messages Spreadhseet					
	SMTP Mess Processed/Hr	SMTP Size Tot KB/Hr	Mess Recip Tot/Hr	Mail Tot KB Trans/hr	Mail Tot Routed SMTP/Hr
Copy these cells--->	0	0	0	0	#VALUE!



Secret

We're into MS Excel for the Analysis

- Next, we open the C:\delme.csv in Excel
 - ◆ Excel knows we want to import it because it's a CSV file
 - ▶ It opens quickly with no further prompts



The Data Is Now in Excel

- The view brought it in sorted by Server and Collection Time
 - ◆ Remember, we'd like to see the number of transactions per hour
 - ◆ With the way this spreadsheet is set up, it's pretty easy to construct a formula where we simply:
 - ▶ Subtract the last hour's number of transactions from this hour's transactions to get the number per hour

A	B	C	D	E	F	G	H	I	J
Server	OS	Collection Time	Day	Hour	User:	Mem Alloc'd	Trans/Min	Trans/Tota	Trans/Hour
USAppDev01	Windows NT 5.0	6/25/2004 16:07	6	16	6	227,821,380	41	1,908,823	
USAppDev01	Windows NT 5.0	6/25/2004 17:07	6	17	5	227,792,466	64	1,912,755	
USAppDev01	Windows NT 5.0	6/25/2004 18:07	6	18	5	227,764,598	50	1,916,463	
USAppDev01	Windows NT 5.0	6/25/2004 19:07	6	19	5	227,744,388	61	1,920,228	

Tricky Calculations – Server Restarts and Stuff

- Except sometimes when servers are restarted
 - ◆ Then the cumulative stats start over
- Or when the next server starts being listed in the statistics
 - ◆ You have to be careful not to subtract without paying attention to these things

A	B	C	D	E	F	G	H	I	J
Server	OS	Collection Time	Day	Hour	User:Mem Alloc'd	Trans/Min	Trans/Total	Trans/Hour	
USAppDev01	Windows NT 5.0	6/26/2004 2:04	7	2	3 234,655,812	58	1,946,368		
USAppDev01	Windows NT 5.0	6/26/2004 3:07	7	3	3 231,761,620	0	1,950,204		
USAppDev01	Windows NT 5.0	6/26/2004 8:08	7	8	3 217,670,374	51	11,149		

A	B	C	D	E	F	G	H	I	J
Server	OS	Collection Time	Day	Hour	User:Mem Alloc'd	Trans/Min	Trans/Total	Trans/Hour	
USAppDev01	Windows NT 5.0	6/28/2004 9:07	2	9	7 220,668,690	197	631,597		
USAppDev01	Windows NT 5.0	6/28/2004 10:08	2	10	12 221,061,694	259	644,588		
USMail01	AIX 1.5	6/21/2004 4:01	2	4	49 840,807,678	305	358,405		

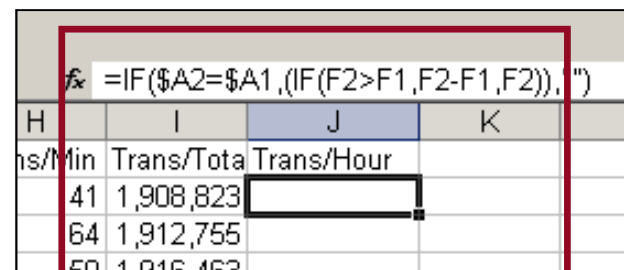
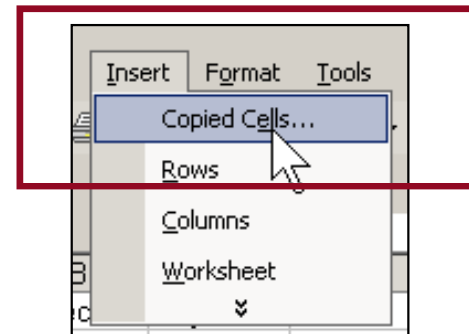
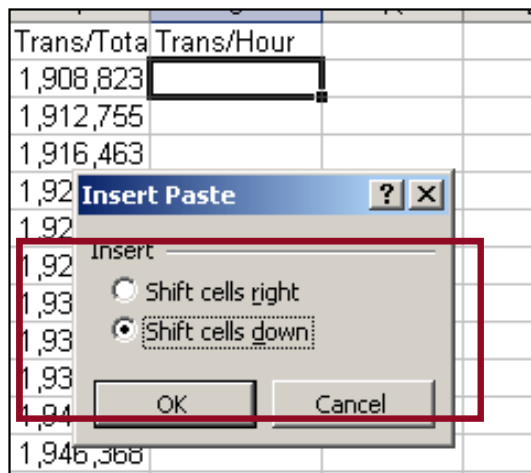
Special Formulas to the Rescue

- To cope with the anomalies in the way the data is listed, I built a few fairly straightforward formulas you can use on your spreadsheets
 - ♦ They are in the master formula spreadsheet
 - ▶ **Just copy it from the cell**

fx =IF(I11=I10,IF(Q11<=Q10,Q11,Q11-Q10),"			
S	T	U	V
	Trnasactions and Users Spreadsheet		
	Trans/Hour		
Copy this cell---->	0		

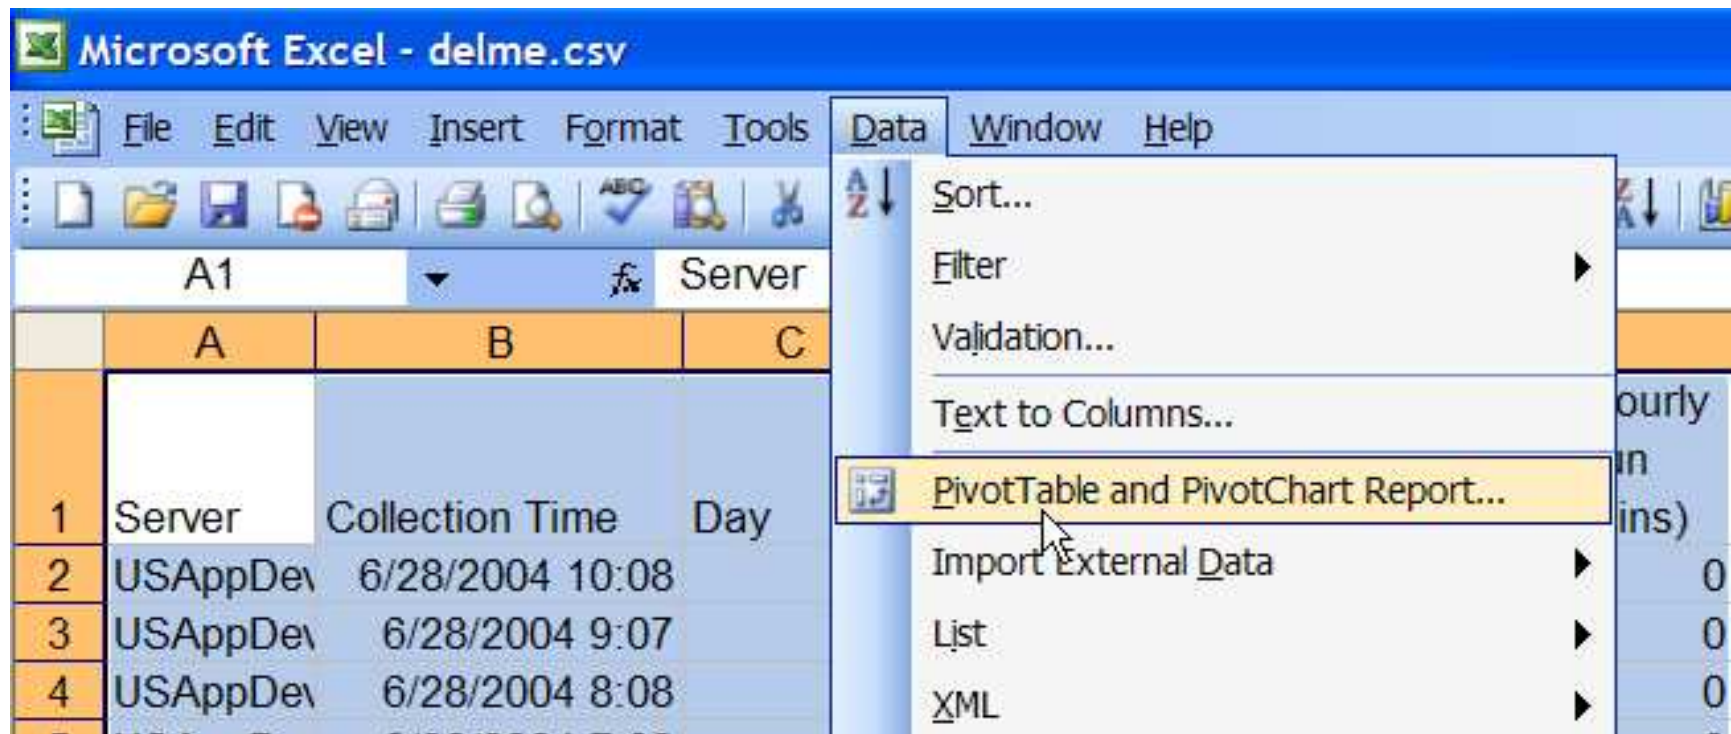
Insert the Copied Cells

- Move to the delme.csv spreadsheet
- Then use the Insert menu to insert the copied cells into your spreadsheet
 - ◆ Move the cells to the right or down to get them out of the way
 - ▶ You'll be copying the proper formula into your spreadsheet
- Copy that formula down your entire column of data
 - ◆ Save your spreadsheet as an XLS



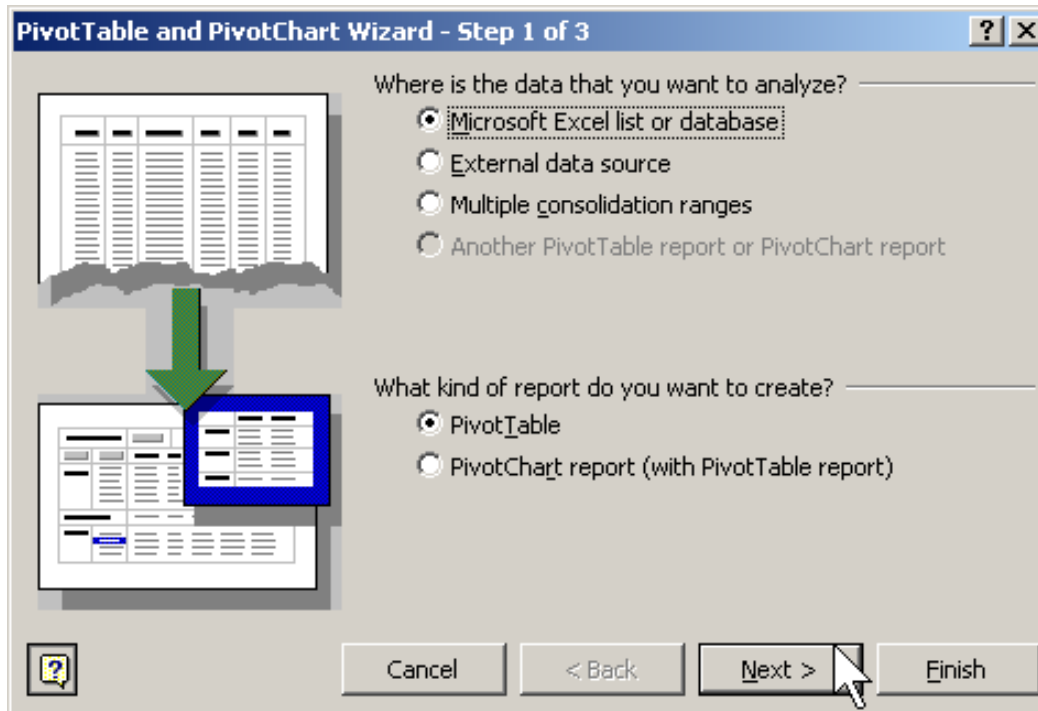
Copy That Cell Down

- We're going to make a Pivot Table with our data
 - ◆ The Pivot Table will take our data and let us easily manipulate it and graph it
 - ▶ **Select all the data, including the column titles, and use the menu to select PivotTable and PivotChart Report**



Take Defaults

- If you're new at this, just take the default answers for the questions Excel asks



**Best
Practice**

The End of the World as You Know It

- It drops you into the Pivot Table function where you have a field list to drag and drop into the table

The screenshot displays an Excel spreadsheet with a PivotTable layout. The layout is defined by blue borders and contains the following text:

- Row 1: Drop Page Fields Here
- Row 3: Drop Column Fields Here
- Column A: Drop Row Fields Here
- Center: Drop Data Items Here

The PivotTable Field List task pane is open on the right side of the screen. It contains the following fields:

- Server
- OS
- Collection Time
- Day
- Hour
- Users
- Mem Alloc'd
- Trans/Min
- Trans/Total
- Trans/Hour

At the bottom of the task pane, there is an "Add To" button and a dropdown menu currently set to "Row Area".

Drag Server to the Column Top

- Drag Server to the column top and Hour to the row names column

The image shows an Excel PivotTable with the following structure:

	Server					
Hour	USAppDev01	USMail01	USMail02	USMail03	USMail04	USMail05
0						
1						
2						
3						
4						
5						
6						
7						
8						
9						
10						
11						
12						
13						
14						
15						
16						
17						

The PivotTable Field List window is open, showing the following fields:

- Server
- OS
- Collection Time
- Day
- Hour
- Users
- Mem Alloc'd
- Trans/Min

The 'Server' and 'Hour' fields are selected in the list. A watermark 'Drop Data Here' is overlaid on the image.

Drag the Data to the Center of the Table

- Drag the data you want to the table itself
 - ◆ It defaults to the “Count of Trans/Hour”
 - ▶ But you’ll want to change it to Average, and format it to look nice, too

The screenshot shows an Excel PivotTable with the following data:

Count of Trans/Hour	Server
0	USAppDev01
1	USAppDev01
2	USAppDev01
3	USAppDev01
4	USAppDev01
5	USAppDev01
6	USAppDev01
7	USAppDev01
8	USAppDev01
9	USAppDev01
10	USAppDev01
11	USAppDev01
12	USAppDev01
13	USAppDev01
14	USAppDev01
15	USAppDev01
16	USAppDev01
17	USAppDev01
18	USAppDev01
19	USAppDev01
20	USAppDev01

The PivotTable Field task pane is open, showing the following settings:

- Source field: Trans/Hour
- Name: Average of Trans/Hour
- Summarize by: Average
- Fields list: Hour, Users, Mem Alloc'd, Trans/Min, Trans/Total, Trans/Hour
- Add To: Row Area

There You Have It

- You now have a nice breakdown of the average number of transactions per hour, per server

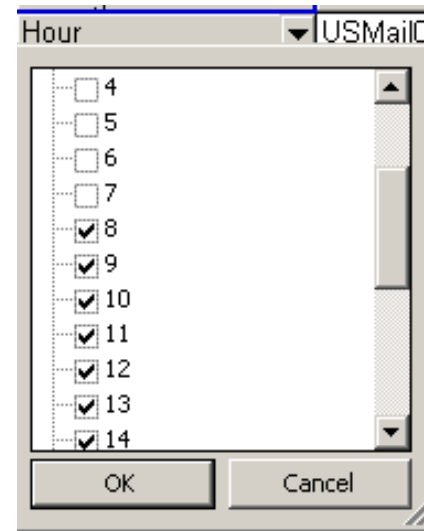



**Key
Feature**

Average of Trans/Hour	Server						
Hour	USMail01	USMail02	USMail03	USMail04	USMail05	USTech01	C
0	15,199	14,524	14,464	27,793	13,939	313	
1	14,324	14,256	14,748	18,637	15,070	376	
2	13,581	13,317	14,019	17,666	15,167	431	
3	12,784	12,654	12,880	16,318	13,064	379	
4	13,753	12,491	12,479	15,746	14,156	389	
5	14,102	12,927	13,237	17,016	14,478	263	
6	14,861	19,453	14,207	17,646	17,247	276	
7	24,999	30,458	49,857	32,147	35,387	268	
8	44,424	56,927	53,572	54,591	64,925	271	

Easy to Manipulate


- It's easy to remove servers and add them back again
 - ◆ And it's easy to pick the hours that you are interested in, too



Average of Trans/Hour	Server						
Hour	USMail01	USMail02	USMail03	USMail04	USMail05	USTech01	Grand Total
8	44,424	56,927	53,572	54,591	64,925	271	47,715
9	64,659	86,529	94,210	98,661	81,473	334	73,359
10	75,279	83,831	102,281	104,936	72,929	360	75,527
11	90,000	89,108	101,054	111,328	74,148	691	76,858
12	68,857	89,255	86,227	113,025	63,418	599	72,203
13	61,046	73,409	65,815	90,267	63,774	406	59,657
14	60,442	74,513	79,908	108,352	62,698	593	63,486
15	111,226	85,224	96,078	138,984	69,284	604	79,699
16	61,215	71,036	91,216	121,049	55,353	632	66,343
17	51,292	75,592	83,396	156,836	53,291	411	72,430
18	27,399	40,396	39,443	95,103	22,676	569	37,726
Grand Total	65,048	75,012	81,495	108,215	62,192	499	65,981

Graphing Your Results

- This is where it really gets cool
 - ◆ Just click on the Chart Wizard
 - ▶ And ...

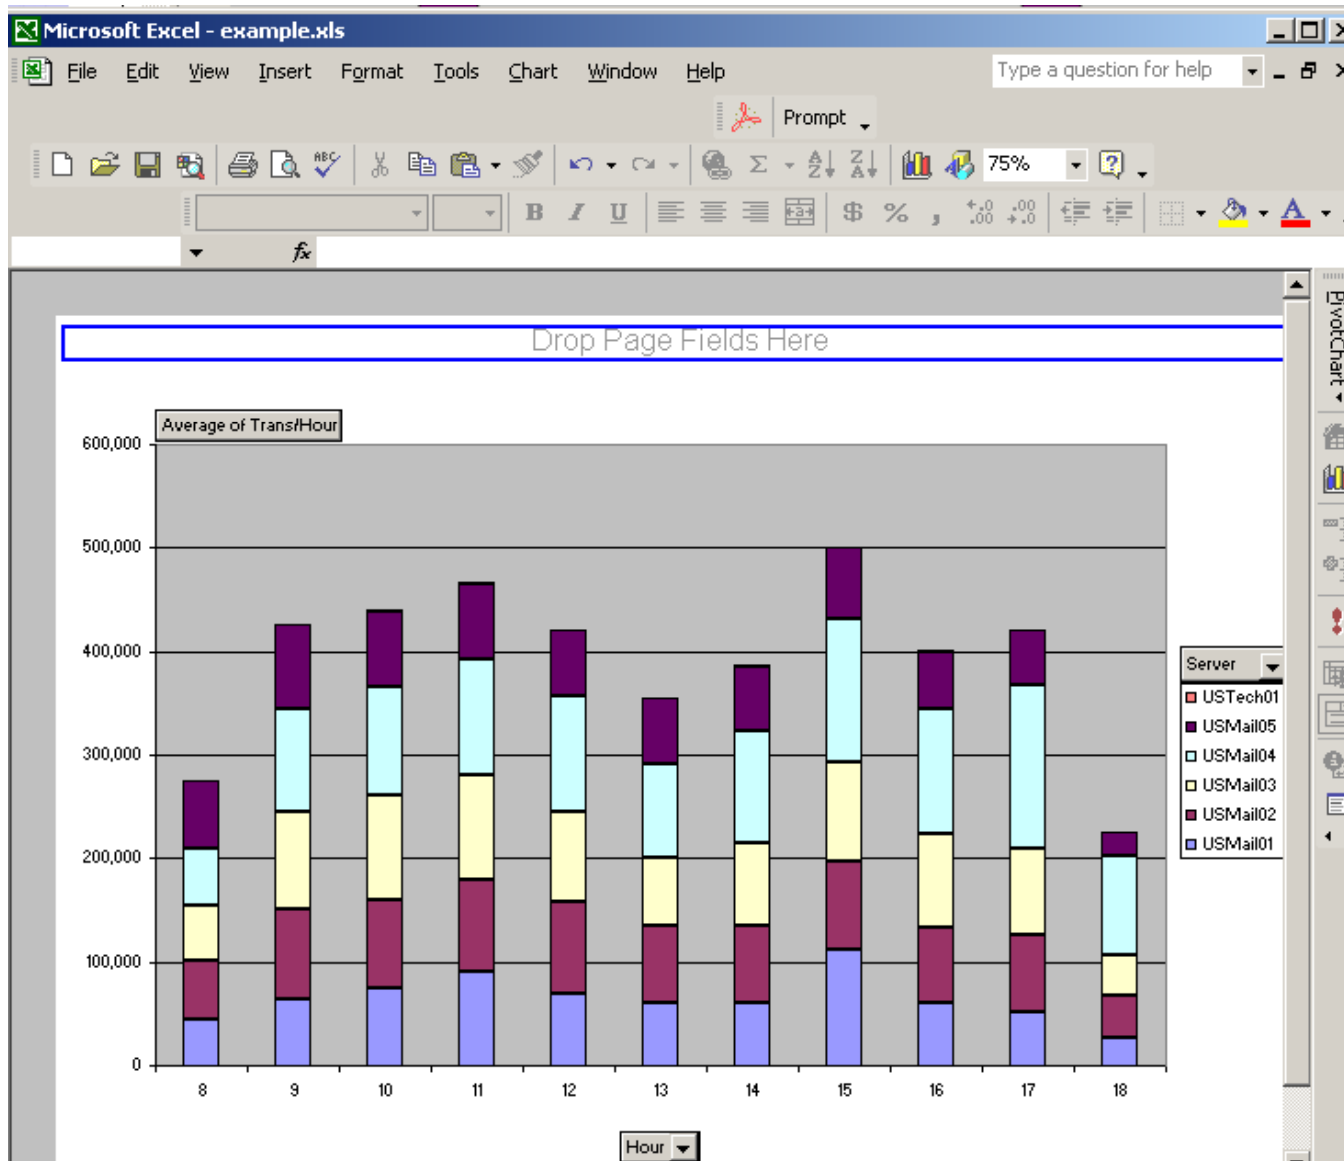


The screenshot shows the Microsoft Excel interface. The menu bar includes File, Edit, View, Insert, Format, Tools, Data, Window, and Help. The toolbar contains various icons for file operations, editing, and data manipulation. A red box highlights the Chart Wizard icon, which is a small bar chart with a mouse cursor pointing to it. Below the toolbar, the worksheet grid is visible. The active cell is C20, and the formula bar shows a function symbol (fx). The worksheet has columns A through H and rows 1 through 5. The data in the grid is as follows:

	A	B	C	D	E	F	G	H
1								
2								
3	Average of Trans/Hour	Server						
4	Hour	USMail01	USMail02	USMail03	USMail04	USMail05	USTech01	Grand Total
5		44 474	56 927	53 572	54 591	64 925	771	47 715

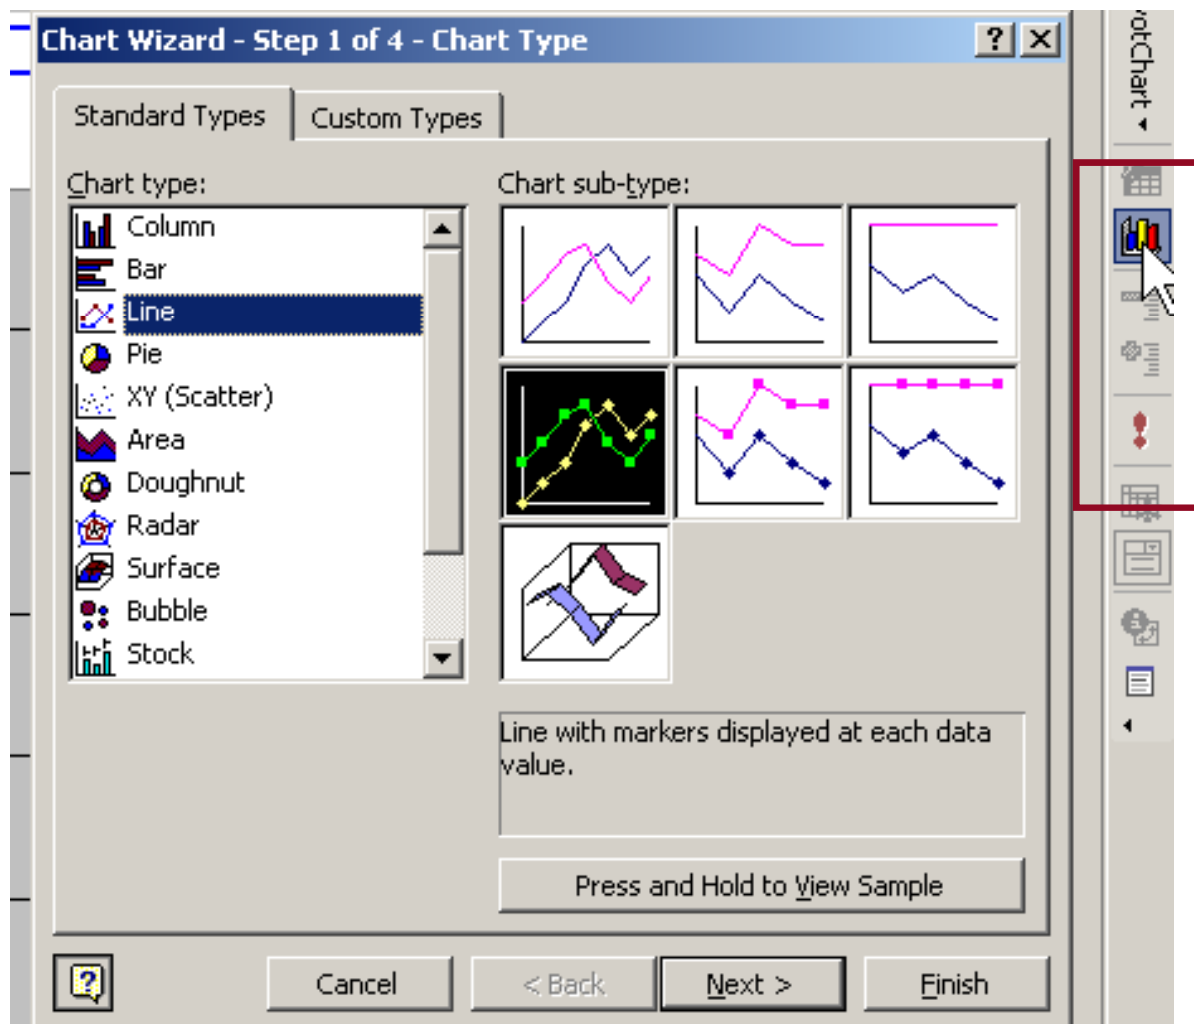
Bingo, You Have an Instant Chart

- Stacked bar isn't what we want, but that was quick!




Line Graph Coming

- Use the icon on the right to change graph types
 - ◆ A line graph is quite effective, most of the time



Here's the Line Graph You Ordered

- Simple, fast, and straightforward
 - ◆ This is an average of transactions per hour



Demonstration: Exporting Data and Creating Pivot Table

Demo



Average Number of Concurrent Users/Hour

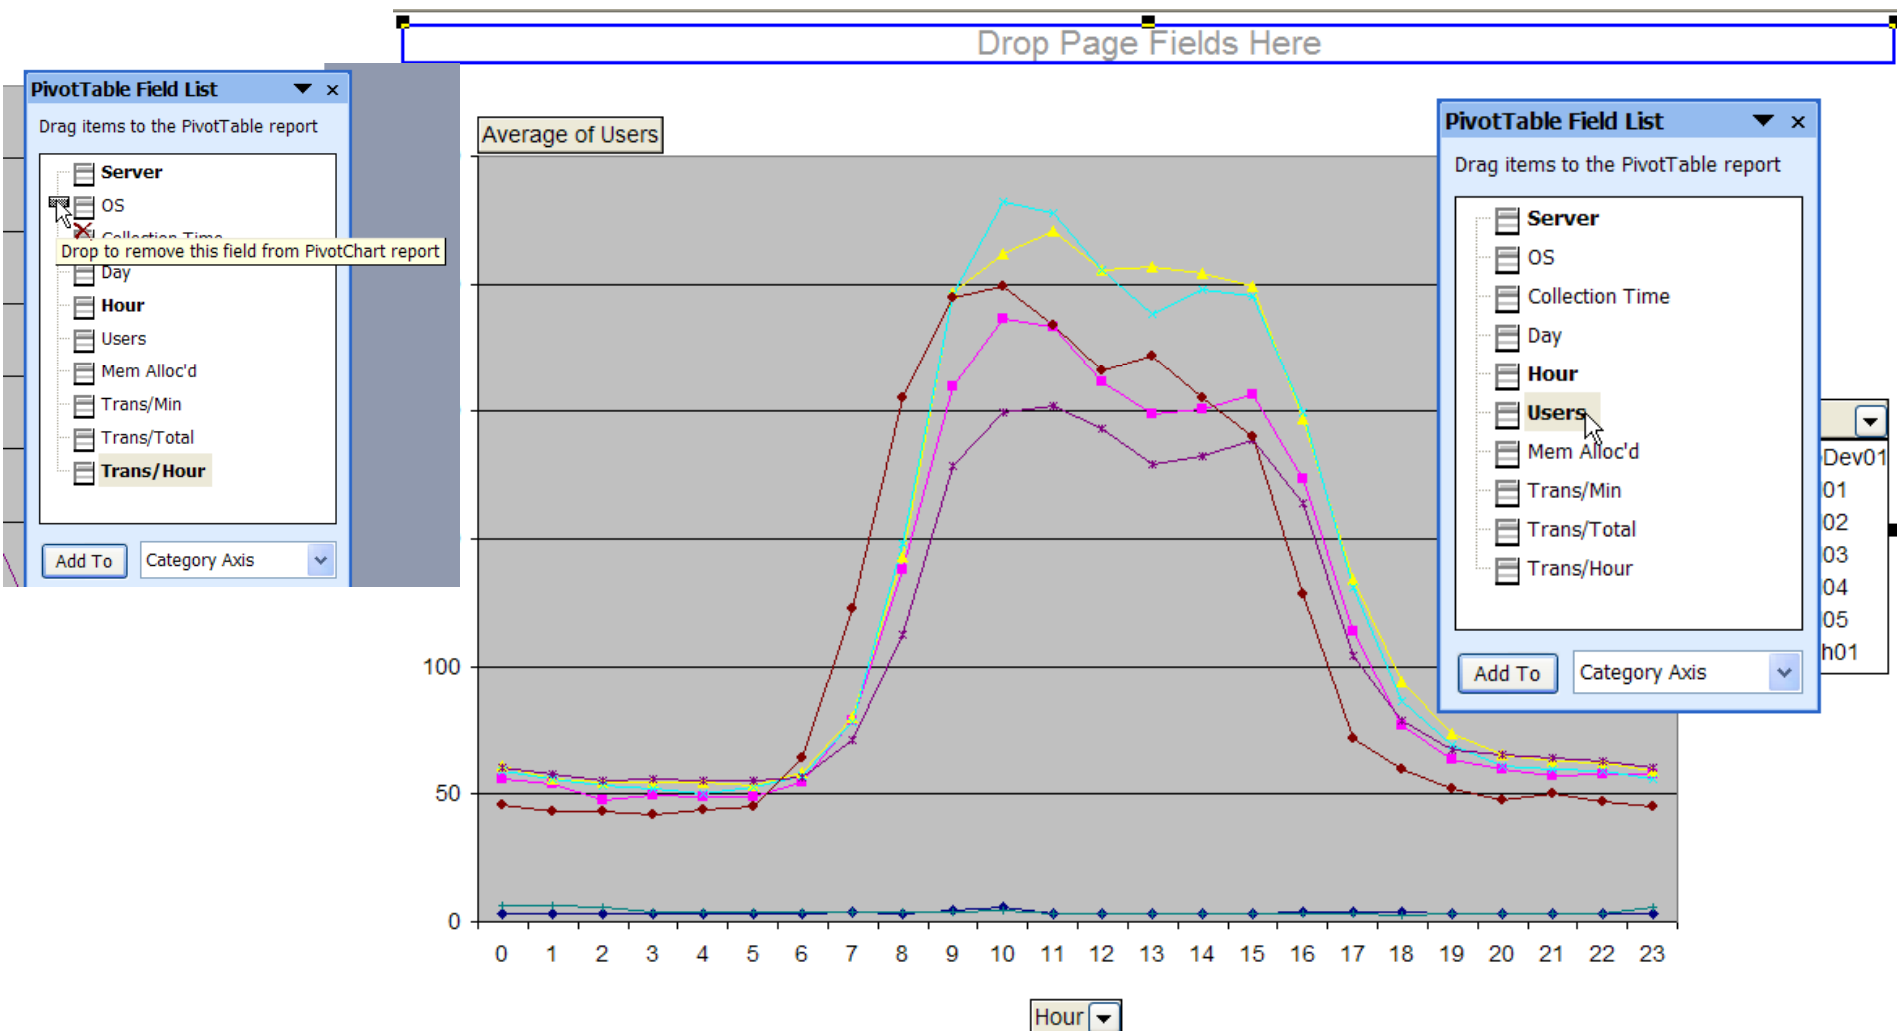
- **This is an extremely valuable statistic**
 - ◆ **Especially when consolidating servers**
 - ▶ **However, there is a Notes.ini variable you must add to servers before this statistic is reliable**
 - ▶ **Here's why ...**
- **When a user connects to a server, they stay connected**
 - ◆ **And are not dropped until they are inactive for four hours**
 - ▶ **This makes it impossible to track actual concurrency because many users may or may not really be active**

Preventing Idle Connections

- To prevent these idle sessions from taking up valuable resources, add this to the Notes.ini of all servers
 - ♦ **Server_Session_Timeout = 30**
 - ▶ **Sets number of minutes of inactivity after which a server automatically terminates network and mobile connections**
 - *Users will not have to re-enter a password if they become active after the time limit*
- The minimum recommended setting is 30-45 minutes
 - ♦ A lower setting may negatively impact server performance
- Now it's easy to chart user concurrency using the same spreadsheet we just developed

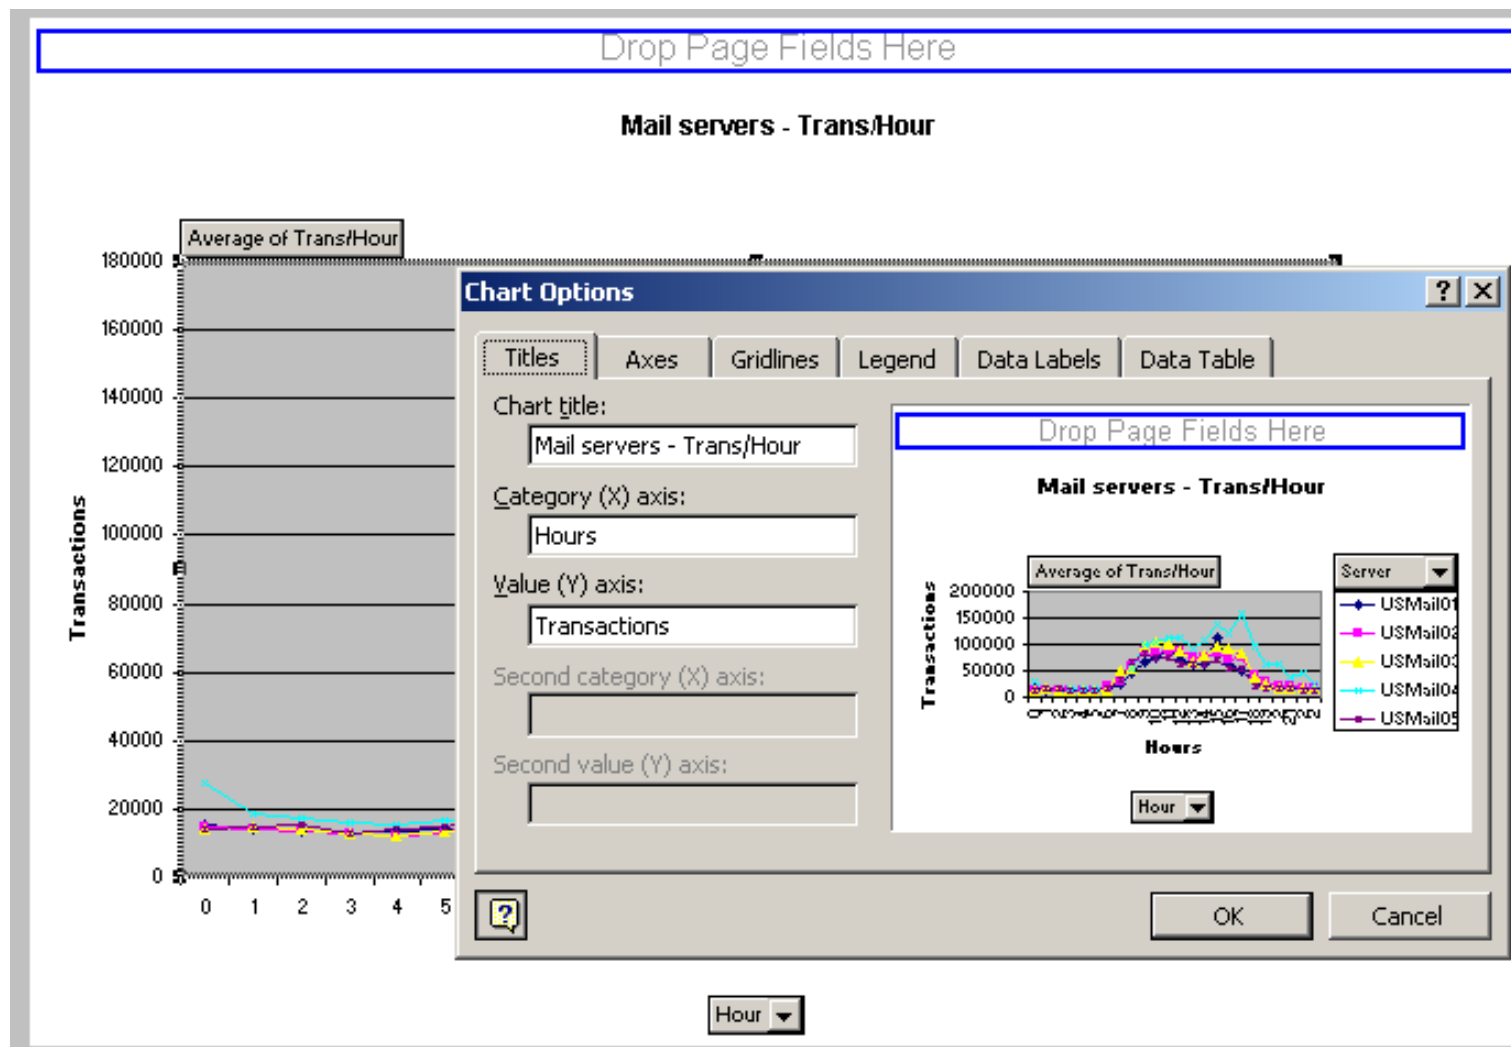
Change the Field List Easily

- It's easy to remove the field Trans/Hour off the chart, and replace it with the Average of Users



Dress It Up for a Presentation

- You can fix it up and format it if you need to make a presentation from the data

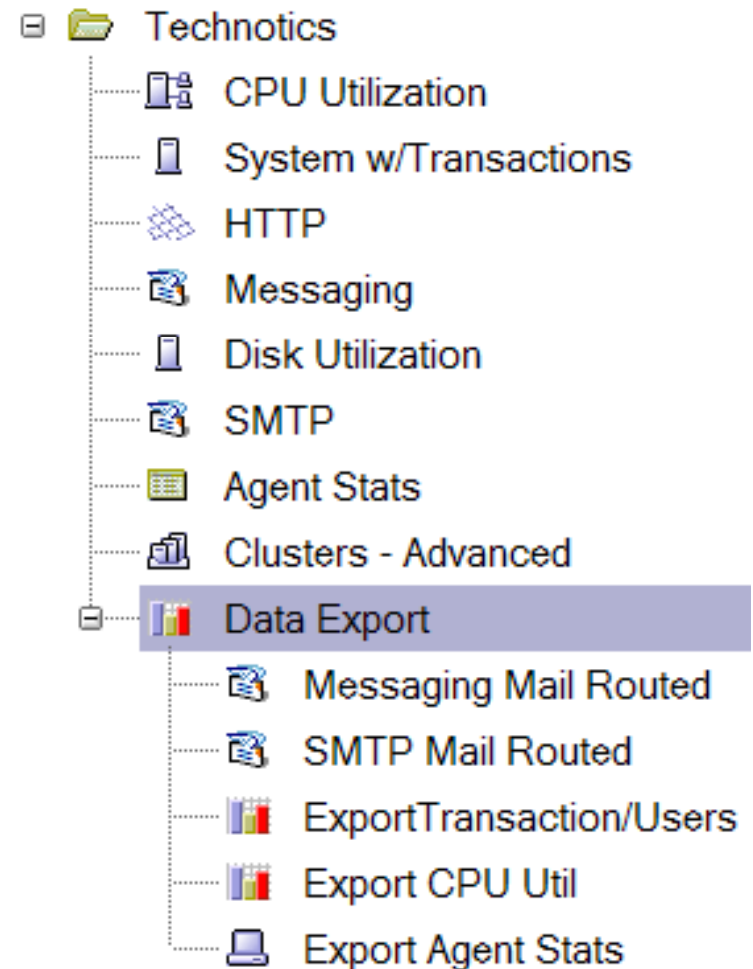


Five Export Views

- There are five different export views on the TechnoticsSTATREP10-9.NSF template from Technotics
 - ◆ Messaging Mail Routed
 - ◆ SMTP Mail Routed
 - ◆ ExportTransaction/Users
 - ◆ Export CPU Util
 - ◆ Export Agent Stats
- Along with the other custom views mentioned earlier



Solution



Messaging Mail Routed and SMTP Mail Routed

- The views for exporting the Messaging Mail Routed and SMTP Mail Routed views use a spreadsheet technique similar to the one used for analyzing transactions per hour
 - ◆ But there are opportunities for analyzing
 - ▶ Average SMTP Messages processed per hour
 - ▶ Average SMTP Message Size processed per hour
 - ▶ Average Message Recipients processed per hour
 - ▶ Average Mail Total Processed per hour

Spreadsheet Concepts Similar


- You will need to copy a group of formula cells instead of just one
 - Insert the copied cells the same way as described earlier in this presentation

SMTP Messages Spreadhseet					
	SMTP Mess Processed/Hr	SMTP Size Tot KB/Hr	Mess Recip Tot/Hr	Mail Tot KB Trans/hr	Mail Tot Routed SMTP/Hr
Copy these cells--->	0	0	0	0	#VALUE!

=IF(\$A3=\$A2,IF(E3>=E2,E3-E2,E3), "")						
I	J	K	L	M	N	
Mail Tot Routed SMTP	SMTP Mess Processe d/Hr	SMTP Size Tot KB/Hr	Mess Recip Tot/Hr	Mail Tot KB Trans/hr	Mail Tot Routed SMTP/Hr	
987						
993	8	51	8	6	6	
999	8	45	8	6	6	
1,005	9	34	9	6	6	


Messaging Mail Routed

- The Messaging Mail Routed export process will allow you to produce a chart like this:




SMTP Mail Routed

- The SMTP Mail Routed will allow you to easily make a chart that looks like this:



Export CPU Utilization

- The Export CPU Utilization will give you a lot of different charts, like this nice one averaging transactions per minute:



What We'll Cover ...

- **Gearing up for advanced statistical analysis**
- **Creating an efficient statistic collection architecture**
- **Customizing the STATREP.NSF (Monitoring Results DB)**
- **Making sure clusters are ready for emergency failover**
- **Mastering the basics of statistical data extraction**
- **Scooping out hidden data to analyze and chart**
- **Wrap-up**

Resources

- **A Domino 6 monitoring and statistics tutorial**
 - ♦ www-128.ibm.com/developerworks/edu/i-dw-ls-dom6stats-i.html
- **Description of HTTP statistics for a Lotus Domino server**
 - ♦ www-1.ibm.com/support/docview.wss?uid=swg21207314
- **Thomas Bahn's blog with code to easily add databases from any and all servers to your desktop**
 - ♦ www.assono.de/blog/d6plinks/Add-Database-From-Servers-To-Workspace
- **Domino performance tuning best practices**
 - ♦ www.redbooks.ibm.com/abstracts/redp4182.html?Open

Things to remember from this presentation

- **The Collect task should only be run on certain servers in the infrastructure that are designated statistic collectors**
 - ♦ **Collect should not be run on every server**
- **Cumulative statistics are only valuable if you perform some very basic calculations**
- **Monitor cluster replication stats weekly to make sure the queues aren't long and the work queue depth is short**
- **High disk queue lengths sometimes indicate failing or misconfigured hardware**
- **There is a Platform.Process.***.PctCpuUtil stat for each process you run**
 - ♦ **Make sure to analyze tasks for their CPU usage**
- **IBM's STATREP database has problems, so don't be shy about creating your own views**
- **Pull exported data into Excel to analyze data in new ways**

Do you have questions about this presentation?



How to contact me:

AndyP@Technotics.com

<http://www.andypedisich.com>

<http://www.technotics.com>