

@HIRSCHOnAPM

JIM HIRSCHAUER

DUMB AND DUMBER

HOW SMART IS YOUR MONITORING DATA



Introduction

- I'm a Performance Geek!!!
- Designed and Implemented Monitoring Architecture for Wachovia Investment Bank and Wells Fargo Managed Services
- I've used many of the enterprise class monitoring tools in existence.
- I currently live, work, and play in Idaho, USA

Right Here!



Agenda



Big Dumb Data



Smart Data Defined



Shifting DR to PR



Smart Data Strategies



Examples




Questions

Big Dumb Data

To quickly identify and remediate the business impact of performance and stability issues.

Why do monitoring tools exist anyway?

What is Business Impact?



Big Data = Enterprise Data Bloating

- Business Data
- Log Files
- Monitoring Data
- Business Intelligence Data
- Legal Data
- Regulatory Compliance Data
- Email
- Etc...

Keep Everything?



Keeping Too Little is Also Bad



Keep Just What You Need



True Story: Oops, that got expensive.

5-7 years ago installed and operated 3 monitoring tools

BTM, APM, and Predictive Analytics
~80 Applications

Ended up with ~50 Management Servers
And 5-10 TB of data

Explore the hidden costs before you decide to implement

The Digital Hoarders are Winning



Gartner Survey

Data Storage



System Performance



Network Bandwidth



40%-60%

False Pretense That Storage is Cheap

- 5 Year Storage Costs: 80% OpEx, 20% CapEx
(2009 IBM Study)
- IT Budgets: Up To 40% Spent on Storage
- \$5-25/GB/month Fully Loaded Cost
 - \$61,440 - \$307,200 Per Year Per TB

Smart Data Defined

Data must be turned into information to be useful.

Heart Rate = 150 bpm

Blood Pressure = 200 over 100

Is the person performing well or not?

Are we talking about this guy?





Or this guy?

Data must be turned into information to be useful.

Eye Color = Brown

Weight = 207 lbs (94 kg)

Is the person performing well or not?

Distance Run = 100 meters

Time = 9.58s



World Record Time=9.69s

Correlation + Analytics Turned Data Into Information




Traditional Monitoring Tools Are Misleading


Core 1 (Hyper-Threading)



Core 2 (Hyper-Threading)



590x333




Resource Spikes **May or May Not** Cause Business Impact

Having a lot of data causes a **false sense of security**.




Your needle is somewhere in there, good luck finding it anytime soon.


We've become addicted to metrics!




What do these charts tell us about application performance or business impact?



(last 6 hours)
CPU % Busy



This is better, but still not good enough.



True Story: Wasted Time.

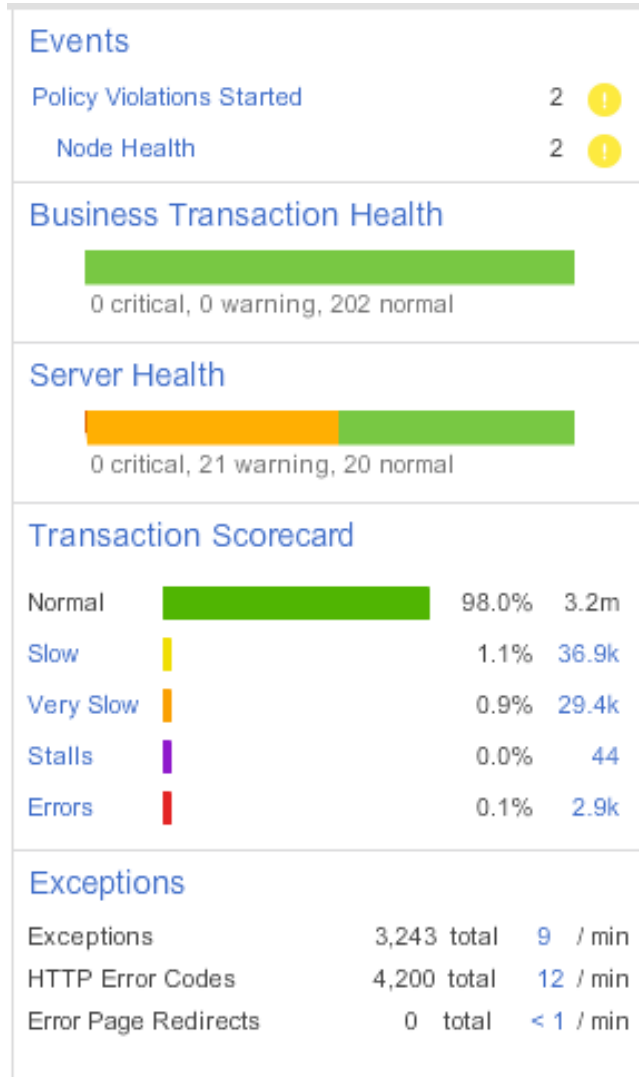
Called onto conf line to help with Sev 1

Confident I had all of the data I needed to figure out the problem

Searched charts for hours

The problem wasn't on my servers in the first place

We need our monitoring platforms to do the heavy lifting for us if we want MTTR < 30 minutes.



Monitor my application from the user AND IT perspective.

Determine what is normal by observation and analytics.

Show me what my application looks like right now using correlation.

Alert me if anything above changes for the worse.

Have the data I need to solve the problem and lead me to the answer quickly.

Disaster Recovery (DR) Needs to Shift to Problem Recovery (PR)



We spend too much time planning for what will probably never happen.



We spend too little time planning for
what happens all too often.


What is Problem Recovery Planning?

PR is a strategy and an organizational mindset.

It's the idea that monitoring is critical to managing applications and ensuring an optimal user experience.

It's the practical implementation of a well defined monitoring architecture.

Monitoring is an afterthought too often.



When a problem occurs...

- Do we have monitoring?
- What kind?
- What are we collecting?
- How long do we have history?

Think about what you need ahead of time.



DB



Network



Log



Infra



App

True Story: Investment Bank Blues

- 40-50 Sev 1 Incidents Per Month
- MTTR ~2 hours
- Executive Mandate to Cut Incidents to Single Digits
- Executive Mandate of 15 Minute or Less MTTR for All Trading Applications

Had It Already

- Infrastructure Monitoring
- NPM – Network Performance Monitoring
- Periodic Database Monitoring

Missing

- APM – Application Performance Monitoring
- Log Monitoring and Analytics
- Always On Database Monitoring
- Predictive Analytics

Added

- APM – Application Performance Monitoring
- Predictive Analytics
- Always On Database Monitoring
- Business/IT Master Dashboard

Significant Results

- Reduced Sev 1s from 45/month to 4/month
- Improved key transaction speeds by 10x
- Reduced MTTR from 3 hrs to 30 mins
- Detected and repaired problems before impact

Cloud Computing is driving the need for PR planning

- Cloud apps are highly distributed so they can take advantage of dynamic scaling
- Highly distributed applications are much harder to troubleshoot
- Use of APM is the fastest way to identify and fix application problems in the cloud

Smart Data Strategies

**MOST APM TOOLS ARE
DUMB**


**MOST INFRASTRUCTURE
MONITORING TOOLS ARE
DUMBER!**

- Single High Traffic Application
- Transmit and store up to 40 TB of monitoring data per year! (Keep Everything)

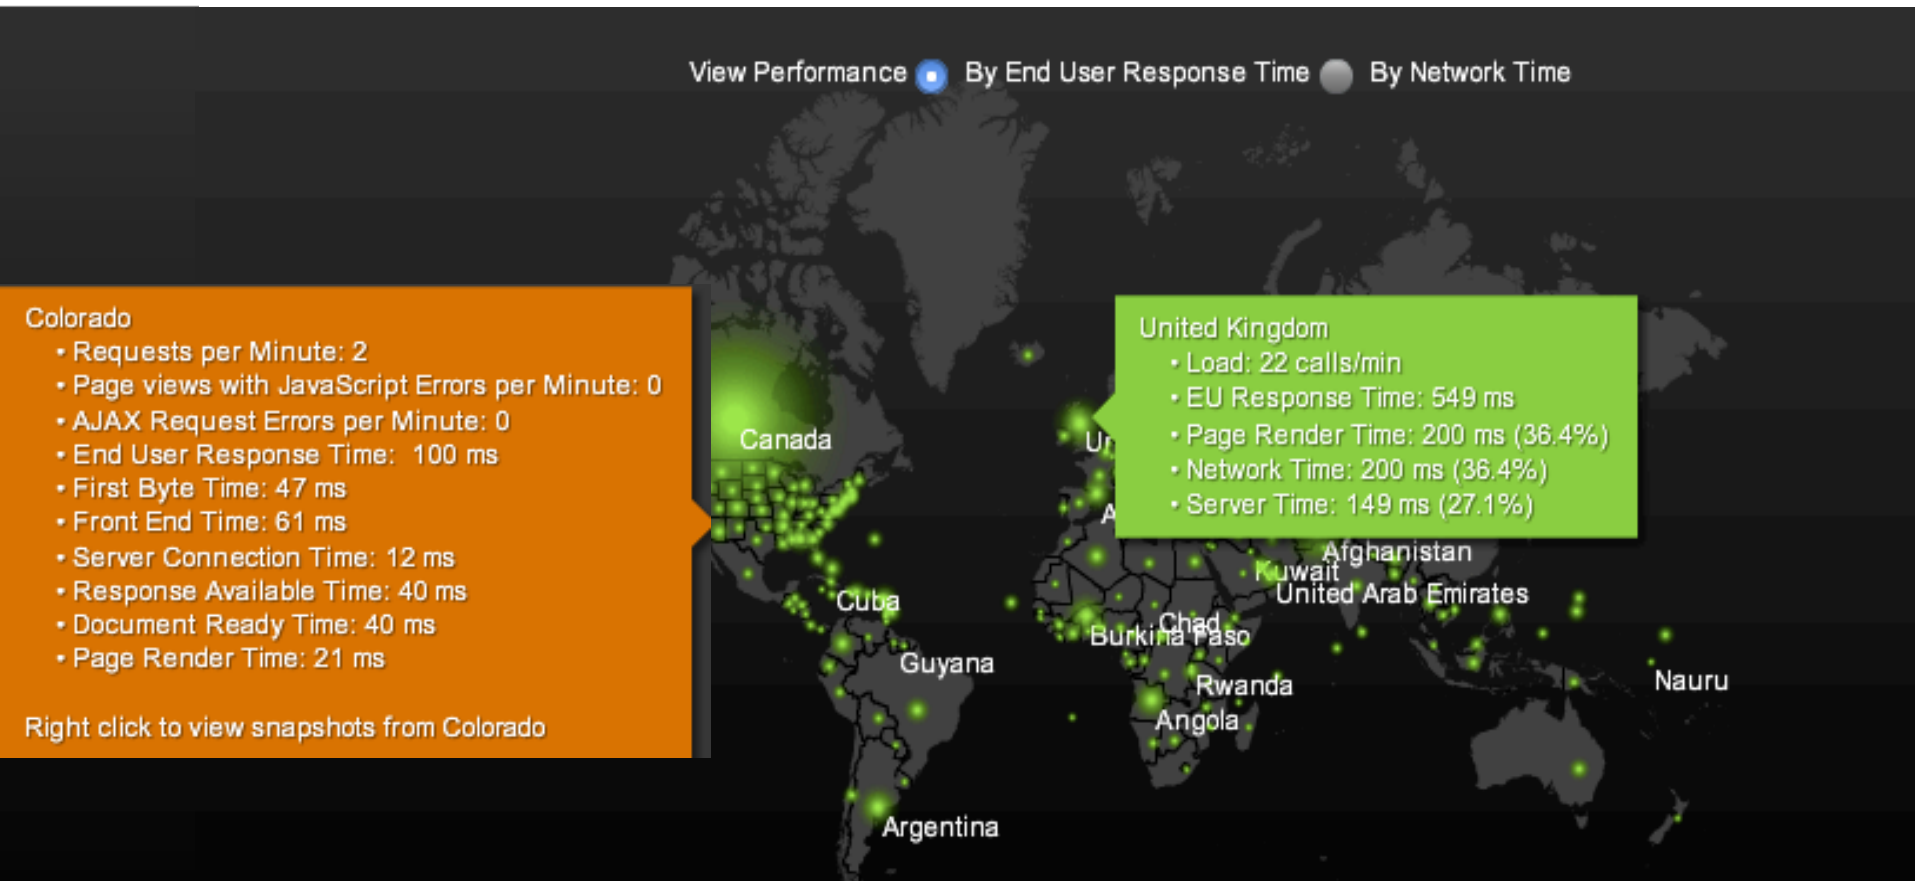
The costs add up.

- Cloud Bandwidth = ~\$5000 per year per application. Charged \$.12 per GB of data out of cloud.
- Storage Costs = \$204,800 per month by end of year 1. Using \$5 per GB per month.
~1.3 Million USD spent at end of 1st year.

We need to save THE RIGHT data






EUE – Key Performance Indicators (KPIs)



EUE – Pages, response time, network time, render time, location performance, etc...

EUE – Key Performance Indicators (KPIs)

d User Experience ▶ Pages & Ajax Requests									
FILTER  Pages  AJAX Requests  iFrames <input checked="" type="checkbox"/> Don't show if i									
	Requests per Minute	Total Number of End	End User Response Time (ms)	Front End Time	Page Render Time	Document Ready	Document Download Time (ms)	Document Processing	First Byte Time (ms)
	2	14	125	47	32	15	0	15	77
	1	9	131	92	12	79	33	46	39
	4	4	24	21	7	15	1	14	2
equies	2	3	11	-	-	-	0	1	10
equies	1	1	8	-	-	-	0	0	8
	1	1	156	148	7	141	1	140	8

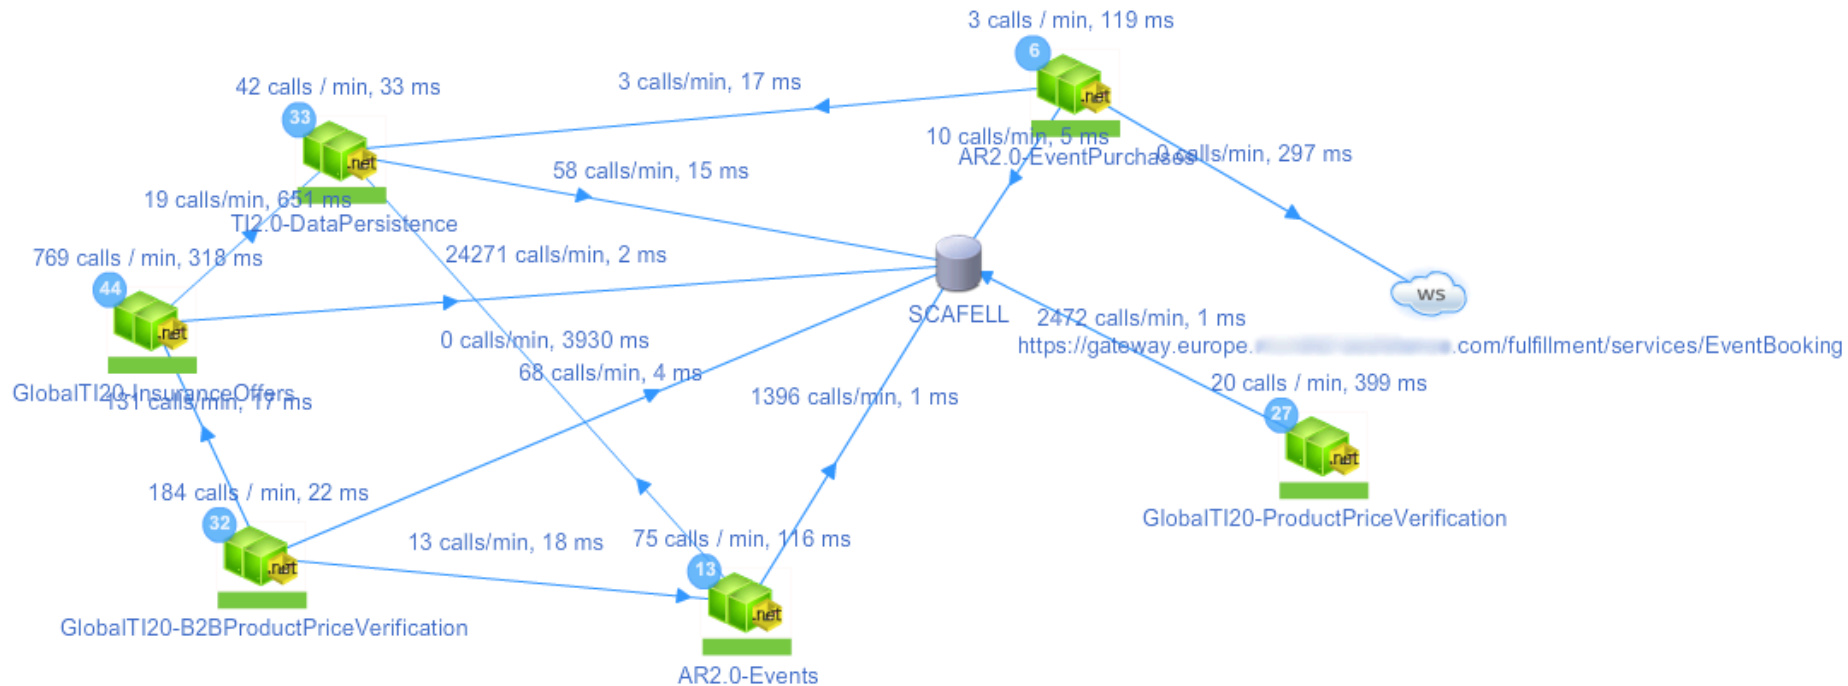
EUE – Pages, response time, network time, render time, location performance, etc...

Business Transaction KPIs

	Name	Health	Server Time (ms)	Max Server	Calls	Calls / min	Errors	Slow Transactions	Very Slow Transactions	Stalled Transactions	CPU Used (ms)	Block Time (ms)	Wait Time (ms)
	All Other Traffic - App		139	22438	39,391	2,626	94	663	332	0	55	0	0
	/ProductDisplay		615	14197	21,661	1,444	8	168	107	0	263	0	0
	/GetColorJSON		410	11208	14,187	946	0	172	190	0	135	0	0
	/ABSwitchView		21	1214	9,981	665	0	62	7	0	15	0	0
	GetMiniCart.execute		1	142	9,408	627	0	0	0	0	0	0	0
	/GetMiniCartHTML.jsp		117	14127	9,406	627	0	121	69	0	49	0	0
	/CategoryDisplay		655	67624	9,114	608	8	42	124	3	260	0	0
	/OrderCalculate		251	15134	7,365	491	1	47	134	0	80	0	0
	/GetUtilityNavHTML		28	10234	7,003	467	0	49	5	0	16	0	0
	OrderItemAdd.execute		191	14864	4,693	313	1	7	83	0	26	0	0
	/ShopRestrict		8	367	4,519	301	0	12	0	0	5	0	0
	/GetMiniWishListHTML		230	23166	4,447	296	0	29	105	0	84	0	0
	/OrderItemDisplayView		375	16407	3,972	265	0	52	50	0	184	0	0
	/GetProductHTML		571	6430	2,634	176	0	31	49	0	214	0	0
	/UserBasicProfileJSONView		24	2664	1,273	85	0	2	2	0	11	0	0
	/Search		2199	30249	1,233	82	0	30	54	0	817	0	0
	/ANFManageShippingInfo		240	16124	1,057	70	0	14	21	0	52	0	0
	/OrderItemDelete		96	15302	959	64	0	2	10	0	7	0	0
	/OrderShippingSectionDisplayView		420	14703	920	61	0	6	3	0	245	0	0
	/ANFShippingInfoUpdate		940	8303	829	55	1	18	21	0	558	0	0
	/ConstantsJS		108	738	695	46	0	9	1	0	84	0	0
	/ShippingMethodDetailsView		121	689	668	45	0	15	0	0	52	0	0
	/GetOrderTotalJSON		45	631	666	44	0	4	0	0	19	0	0

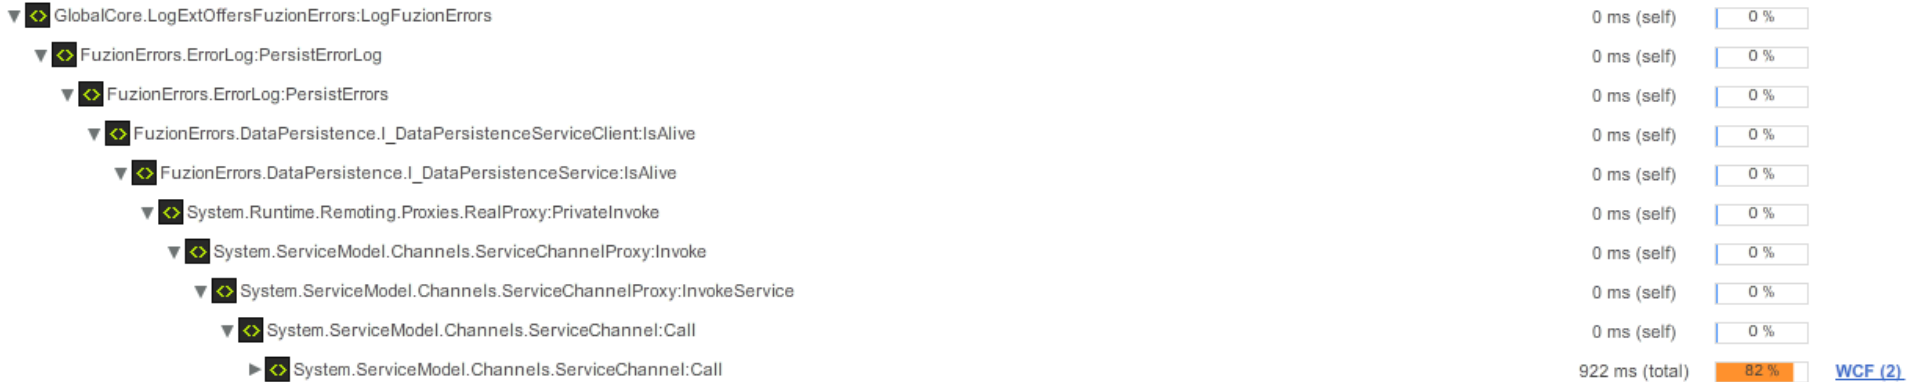
BTs – Response time, count, rate, errors, CPU Used, CPU Block, CPU Wait, etc...

Application Flow KPIs



Application Flow – Active nodes, active tiers, node response time, tier response time, **external service response times**, etc...

Deep Diagnostics – We don't need to save these forever.



This screen displays all of the method calls in the call graph sorted by time

Name	Method Time (ms)	External Calls
System.ServiceModel.Dispatcher.RequestChannelBinder:Request	922 ms (self) 82 %	WCF (2)
System.Data.Common.DbCommand:System.Data.IDbCommand.ExecuteReader	31 ms (self) 2.8 %	ADO.NET
Custom Entry Point - MA.Gateway.WebServices.TravelInsurance.ProductPriceServiceExt:GetProductPriceOffers	31 ms (self) 2.8 %	
System.Data.Common.DbCommand:System.Data.IDbCommand.ExecuteReader	31 ms (self) 2.8 %	ADO.NET (2)
System.Data.OracleClient.OracleCommand:ExecuteNonQuery	31 ms (self) 2.8 %	ADO.NET
System.Array:Copy	16 ms (self) 1.4 %	
System.Data.Common.UnsafeNativeMethods:OCIStmtFetch	16 ms (self) 1.4 %	
GlobalInsuranceNucleus.DomainModel.PartnerRelatedObjects+<>c__DisplayClass4:<GetOfferPackMessagesById>b__0	16 ms (self) 1.4 %	
Castle.DynamicProxy.Invocation.AbstractInvocation:ctor	16 ms (self) 1.4 %	
System.Decimal:Compare	15 ms (self) 1.3 %	

Don't be this guy...



Plan ahead, anticipate your needs, keep your organization nimble, powerful and purpose built.




Example



- Video Streaming
- AWS Deployment
- Highly dynamic environment
- ~10,000 JVM Nodes
- Doing it right

Netflix

The Netflix logo, consisting of the word "NETFLIX" in white, bold, sans-serif capital letters on a red rectangular background.

Collecting over 1 million metrics per minute.

What's the point(s)?

- Big data isn't a bad thing as long as it is serving a purpose.
- Big monitoring data slows down MTTR and drives up both OpEx and CapEx.
- Focusing on Problem Recovery will help you figure out your architecture, tools, and process.
- Don't be a digital hoarder!!!

Questions???

Thank You

