

What Takes Us Down?

What are the causes of IT service disruption? With access to an e-mail archive recording both planned and unplanned events, I figured I could identify ways to reduce downtime. This turned out not to be as easy nor as useful as I had hoped: the exercise raised questions but little that was actionable. Still, the path I took may help you analyze your own data.

Environment

I work at the Fred Hutchinson Cancer Research Center, a non-profit biomedical research institute specializing in cancer and infectious diseases. I pay attention to *deep infrastructure*: power, cooling, cabling, transport (Ethernet, IP, WiFi, Fibre Channel), interstitial services (DNS, DHCP, authentication, directory services), e-mail, storage, file services. These days, I spend my time managing our Problem Management process and leading Root Cause Analysis efforts. (Yes, the ITIL borg is extending its filaments into our brains.)

In the mid-90s, the network team started posting service-affecting incidents, both planned and unplanned, to an e-mail list. Over time, more and more departments followed suit, and more and more techs subscribed to the list. We've refreshed that list server over the years, trashing its archives each time. The current archive starts October 2000.

In our culture, Planned downtime goes over well – we negotiate Service Level Agreements (SLAs) with our divisions specifying when we can take down applications; we notify users; they modify their work-flows to dodge the windows during which we are disrupting service; everyone is happy (for some value of happiness). But Unplanned downtime is another matter – no one enjoys that, and we invest effort to avoid it.

Today, the Center employs 2500 staff with an annual budget of \$450 million, 85% from federal grants and contracts. 8000 active Ethernet ports, 11,000 active IP addresses, 450KW data center cooling, 1PB+ mass storage, national and international collaborations. Roughly 30% of the end-stations run Windows, another 10% Linux, 5% OS X, and the remainder fall into the miscellaneous category (IP phones, printers, etc ... well most of those run Linux, but I want the Linux figure to reflect desktops/laptops/servers only).

The Outages List

Techs send e-mail to this list using a standardized format. In our lingo, all service-affecting events are called *Outages*.

Subject: Exchange 2003 Cluster Issues
Severity: Critical (Unplanned)
Start: Monday, May 7, 2012, 11:58
End: Monday, May 7, 2012, 12:38
Duration: 40 minutes
Scope: Exchange 2003
Description: The HTTPS service on the Exchange cluster crashed, triggering a cluster failover.
User Impact: During this period, all Exchange users were unable to access e-mail. Zimbra users were unaffected.
Technician: [xxx]

Figure 1: Example of an Unplanned Outage

Subject: H Building Switch Upgrades
Severity: Major (Planned)
Start: Saturday, June 16, 2012, 06:00
End: Saturday, June 16, 2012, 16:00
Duration: 10 hours
Scope: H2 Transport
Description: Currently, Catalyst 4006s provide 10/100 Ethernet to end-stations. We will replace these with newer Catalyst 4510s.
User Impact: All users on H2 will be isolated from the network during this work. Afterward, they will have gigabit connectivity.
Technician: [xxx]

Figure 2: Example of a Planned Outage

You can't see this in the template above, but we also categorize Outages by *Window*:

Prime	Monday - Friday 7am - 6pm
SLA	Sunday 8pm - Monday 4am or Wednesday midnight - 4am
Shoulder	Any other time

Figure 3: Windows

Yes, that Service Level Agreement (SLA) window looks pretty darn generous ... we can take down services every single week for eight hours straight?! Conceptually, yes, but in practice, research institutes live and die by grant applications. Those grant applications have submission deadlines, and those deadlines pop up almost every week. Thus, most of those SLA Windows get blocked.

The Severity field has intention glued onto it via parentheses: Planned (we intended the Outage) or Unplanned (we were surprised). Severity itself can contain the following values:

Drama	Most of the Center for 60+ minutes during Prime time
Critical	One or more buildings or divisions
Major	Multiple floors or multiple departments
Minor	One floor or one department
None	No end-user effect

Figure 4: Severities

You might ask why we bother to have a Severity of None – doesn't sound like an Outage if the end-user impact was None, right? Well, the motivation is two-fold. First, most of us want to be kept informed of changes to the environment (because what you change might in fact interfere

with something I do), and the Outages list serves as that forum. Second, we're trying to flush out errors in our understanding of the environment: if I claim that an event did not cause a service disruption and you know differently, you'll tell me (and then I'll send a correction to Outages).

Complicated and subjective? Yup.

The Outages Database

So I figured I'd write code to grab the list archives and crawl through them, creating a database entry for each Outage. How hard could this be? After all, we use this structured template ... Ahh, the naiveté and eternal optimism of youth: two weeks and 2500 lines of Perl later (the grossest code I've ever written), I ran it, took the partially processed results, imported into my database, and started scrubbing manually. Turns out we don't follow the template exactly: I never knew there were so many ways to write a date/time stamp, techs twink with the spelling of key words regularly, techs tend to send multiple messages describing each Outage (the first announcing the event, the last announcing the completion of the event, and often others in between correcting errors or adding new information) ... I'll quit whining here.

Furthermore, I wanted a feel for Service and Cause, neither of which is specified in the template. I added those during my manual passes.

<u>Service</u>	<u>Description</u>
Application	End-user facing apps (minus e-mail, MIS, and printing)
Email	Exchange, Zimbra, mail relay, spam/malware scrubbers
HPC	High Performance Computing
Interstitial	DHCP, DNS, NTP, authentication, directory services
MIS	Financial Management / Human Resources systems
Power	Electricity
Print	Print servers
Storage	Anything providing file or block services
Transport	Ethernet, WiFi, Remote Access, Fibre Channel
Virtualization	VMWare, Xen
Voice	Telephones and pagers

Figure 5: Services

This list reflects the focus of the groups who post to Outages. Most of the Application support groups do not post; I lump their contributions into a single category.

More interestingly, I wanted to identify the Proximate Cause of each Outage – again, not something defined in the template, so I added this during my manual passes, interpreting the Description field, dusting off memories, and making a judgment call.

<u>Cause</u>	<u>Description</u>
Cockpit Error	Techie mistake (fat finger, config error, bump the power cord)
Design	Service didn't behave as the designer intended
External Services	Service provider issue (electric utility, ISP, telecom carrier)
Hardware Failure	The magic smoke escaped
Maintenance	Patching, database compression, shuffling data, minor fixes
Malware	Virus or worm infection
Overload	Too much of a good thing
Software Bug	Memory leak, unhandled exception, Blue Screen of Death
Testing	Validating expected behavior, typically involving high-availability
Unknown	Never figured it out
Upgrade	Adding major functionality (new gear, major software update)

Figure 6: Causes

There's a lot of fuzz here. When we popped a ceiling tile trying to trace cables and knocked an unsecured electrical wire loose, this triggered an Emergency Power Off (EPO) event in our largest data center. I categorized the Service as *Power* and the Cause as *Design Failure*. I could have categorized the Service as *Application* (every application in that data center went down) and the Cause as *Cockpit Error* (the electrician who installed the EPO circuits intended to screw that wire into place but forgot). I didn't because I try to push Cause down the OSI stack (Power sits a whole lot lower than Application), and I try to pick the Proximate Cause as opposed to the Root Cause: at the moment we popped the ceiling tile, Power did not behave according to the data center Designer's intent.

Or, to take another example, if the server stayed up, but became too slow to be usable on account of too many users, I categorized that as *Overload*. If the server crashed because it ran out of RAM (on account of too many users), I categorized that as *Software Bug*.

Yup, pretty darn subjective.

Notes

- External Services isn't really a Cause ... but since the Service Provider space is opaque to us (did the WAN circuit go down due to human error, an unannounced Upgrade, or a Power event?), we lump them together.
- Hardware Failure lumps together both Unplanned events, during which the magic smoke escaped, and Planned events, during which we replace, say, a disk controller which is failing diagnostics but hasn't actually fried yet.
- Maintenance is driven by patching, mostly Windows patching.
- Testing is driven by the network team, which reboots redundant switches and routers monthly.

Results

I ended up with ~2300 Outages [1] spanning the last 11+ years. Is that an undercount? Definitely. Departments vary in how frequently they report: some use Outages rigorously,

others not at all. And of the entries in Outages, I threw away hundreds which my code didn't parse or which were too terse for me to categorize manually.

Q: How often are we surprised?

A: We're surprised half the time

Planned: 55%, Unplanned 45%

Q: What takes us down?

A: Software Bugs take us down

For Unplanned Outages, Software Bugs are the dominant contributor. And for Planned Outages, a third arise from Maintenance, which is driven by patching, i.e. fixing Software Bugs.

Planned		Unplanned	
<u>Cause</u>	<u>Proportion</u>	<u>Cause</u>	<u>Proportion</u>
External Services	2%	Cockpit Error	13%
Maintenance	32%	External Services	7%
Other	14%	Hardware Failure	12%
Testing	11%	Other	7%
Upgrade	41%	Software Bug	61%

Figure 7: Planned vs Unplanned by Cause

Q: When do we go down?

A: In the middle of the day

If Unplanned Outages were to occur at random, regardless of time of day, we would predict that some Unplanned Outages would land during the Prime window (55 hours/week), most during the Shoulder window (101 hours/week), and a few during the SLA window (12 hours/week). But in fact, we see that far more land during Prime time than we would expect based purely on chance – perhaps because our users are exercising the systems and uncovering bugs in the process.

Window	Predicted	Measured
Prime	33%	67%
Shoulder	60%	31%
SLA	7%	2%

Figure 8: Unplanned Outages by Window: Predicted vs Measured

Q: What Causes induce the most pain?

A: The same causes which induce major and minor pain

I tried slicing and dicing in other ways but did not uncover new information. For example, when focusing just on the most painful events (Drama and Critical), Causes break down pretty much the same as they did when considering all Severities:

Planned		Unplanned	
<u>Cause</u>	<u>Proportion</u>	<u>Cause</u>	<u>Proportion</u>
External Services	4%	Cockpit Error	17%
Maintenance	30%	External Services	9%
Other	12%	Hardware Failure	8%
Testing	16%	Other	11%
Upgrade	38%	Software Bug	55%

Figure 9: Planned vs Unplanned by Severity (Drama + Critical Only)

Q: How often does it hurt a lot?

A: Severity Shows a Normal Distribution

We experience a few of the really painful Drama Outages and a few Outages with no end-user effect: most land in the middle.

	<u>Planned</u>	<u>Unplanned</u>
Drama	0%	5%
Critical	16%	19%
Major	46%	35%
Minor	34%	39%
None	4%	2%

Figure 10: Planned vs Unplanned by Severity

Q: What breaks most often?

A: Transport and Email are weak spots

But see caveats below.

	<u>Planned</u>	<u>Unplanned</u>
Application	18%	16%
Email	20%	21%
Other	16%	20%
Storage	14%	8%
Transport	32%	35%

Figure 11: Planned vs Unplanned by Service

Reality Check

Fuzzy Data

Those cute tables with numbers in them look good ... but as I apply cultural knowledge, I lose confidence. For example, the network team founded the Outages list; the e-mail team jumped onto the bandwagon shortly thereafter: these two groups have been posting the longest and have become ruthless about reporting every event, no matter how embarrassing. Furthermore, they have the most mature monitoring systems, reporting even minor hiccups. Are Transport and Email our most fragile services? Or do they top the list because of cultural factors: habit, conscientiousness, visibility?

Cockpit Error

Reading thousands of descriptions of Outages gave me a chance to smile – I remember many of these events from personal involvement and know each of the techs posting to the list. Senior techs tend to acknowledge their errors directly, using language like *I fumbled the configuration ... I accidentally typed rm -rf * from root... I broke the Internet connection...* whereas junior techs tend to slide into passive voice and circuitous language when they describe their errors: *The service went down during trouble-shooting... It was discovered that the configuration file contained an error... On investigation and after analysis, the power cord was found to be detached...* Where possible, I flagged the Cause as Cockpit Error, but I'm confident that I missed plenty. For that matter, I suspect that Cockpit Error leads to unreported Outages, as techs try paddling up the Nile in their efforts to dodge embarrassment. We have a remarkably shame/blame free environment – as far as I know, no tech has ever been fired for making a mistake and causing an Outage. In fact, management likes to stress that making mistakes is how we learn (yeah, OK, sometimes they look a little nervous when they make this point, but still, the sentiment is there). How can we boost the Cockpit Error reporting rate?

Who else quantifies this stuff?

A casual search turned up a handful of studies in this space, with varying sized data sets (typically 100 – 1000 incidents spanning 1-5 years).

	Gray [2]	Kuhn [3]	Enriquez [4]	Oppenheimer [5]			Offord [6]	Kendrick
Published	1990	1997	2002	2003			2011	2012
				SP1	SP2	SP3		
Cockpit Error	13%	25%	38%	33%	36%	19%	42%	17%
Software	58%	14%	7%	27%	25%	24%	38%	55%
Hardware	18%	19%	30%	25%	4%	10%	-	8%
Other	11%	42%	25%	10%	31%	33%	20%	20%

SP = Service Provider

Figure 12: Similar Surveys

I'm skeptical that I'm comparing apples to apples here – both environments and methodologies vary widely. For example, Gray, Kuhn, and Enriquez were all analyzing data sets taken from homogenous systems (Tandem Computers and the Public Switched Telephone Network), while Oppenheimer, Offord, and myself are analyzing heterogeneous environments: Windows/Linux-based systems running on IP/Fibre Channel networks. Or, to take another example, Offord extracts his data set from the log of Root Cause Analysis jobs his company has performed for customers – not exactly Outages but rather long-running Problems. In 42% of their cases, the Problem was fixed by making a Configuration Change, which I recategorized as Cockpit Error, in order to fit his data into my taxonomy – probably not a precise match.

Tentatively, I see all these data sets directing our attention toward software flaws and operator fumbles as places for improvement.

What to Do?

Software Bugs

For us, our Unplanned downtime is driven by Software Bugs (~60%). We know that we lag on patching. When a service fails repeatedly, we'll investigate and often find a patch addressing the issue which the vendor shipped months or years prior. I would like to think that if we patched more regularly, we would convert Unplanned Outages into Planned Outages. Still, this is a tricky area – most of our teams don't have test environments (we are non-profit after all) – so we test patches by running them in production, and as we all know, patches can fix issues we weren't having while introducing new issues. How many Unplanned Outages would we dodge by patching more aggressively?

Testing

Until recently, the network group tested their redundant routers and switches monthly, rebooting them in series, analyzing failure, fixing the issues they uncovered (typically Cockpit Errors, e.g. misconfigurations), working with sys admins to fix misconfigured servers (servers which weren't configured to take advantage of the dual Ethernet switches in data centers) [6], helping the security groups buff up highly available firewalls. Of our really painful Planned Outages, Testing contributed 16%. I would like to think this approach saved us a similar number of really painful Unplanned Outages and thus was a win. On the other hand, Testing requires substantial staff time. How to quantify the costs and benefits?

Insights

I have been struck by the number of axes on which one can measure an incident. Each of the authors I cite developed their own taxonomy. To recap, here's mine:

Function	Our Term	Description
Pain Level	Severity	Drama, Critical, Major, Minor, None
Intention	Planned	Planned or Unplanned
Time Frame	Window	Prime, Shoulder, SLA
End-User Impact	Service	The thing that went down
Proximate Cause	Cause	What caused the downtime

Figure 13: Taxonomy

I am troubled by how subjective my categorization process is – I made multiple passes through the database, recategorizing as I became more familiar with my data – nevertheless, I expect that I made inconsistent choices. Also, many Outages don't fit the taxonomy cleanly: what to do with a Planned Outage which incurred Unplanned consequences? Or an Outage which knocked out multiple Services? And Cause remains tricky – an Outage has so many causes, how to pick just one?

Still, at the end of the day, I'm headed back to Problem Management meetings to suggest Patching and Testing as ways to convert Unplanned events into Planned ones.

Doubt is uncomfortable; certainty is absurd. --Voltaire

[1] See <http://www.skendric.com/problem/incident-analysis> for the summarized data.

[2] Jim Gray, "A census of Tandem system availability between 1985 and 1990", *IEEE Transactions on Reliability*, vol. 39, no. 4, pp. 409-418, Oct. 1990. On p.6, I used the All Faults column and categorized maintenance + operations + process as Cockpit Error.

[3] Richard Kuhn, "Sources of Failure in the Public Switched Telephone Network," *IEEE Computer*, vol. 30, no. 4, pp. 31-36, April 1997, doi:10.1109/2.585151. I counted only errors from telco staff as Cockpit Error, allocating 'Human error – external' to Other.

[4] P. Enriquez, A.B. Brown, and D. Patterson, "Lessons from the PSTN for Dependable Computing," *Proceedings 2002 Workshop Self-Healing, Adaptive, and Self-MANaged Systems (Shaman)*, 2002, pp. 1–7. Again, I allocated 'Human error – external' to Other.

[5] David L. Oppenheimer, A. Ganapathi, D. Patterson, "Why Do Internet Services Fail, and What Can Be Done About It?" USENIX Symposium on Internet Technologies and Systems 2003.

[5] Paul Offord. "RPR Statistics". Advance7, October 2011. Web. October 2011. I map the sum of *Bug Fix* and *Programming* into my Software Bug. *Programming* is Advance7's term for a bug fixable by internal resources, e.g. a bug found in an in-house application, while *Bug Fix* is Advance7's term for a bug found in software acquired externally, e.g. commercial or open-source.

[6] Stuart Kendrick. "Testing the Transport Side of Highly-Available Hosts". *login* vol 36, no. 3, pp. 33-43 June 2011.