

s11n - an Object Serialization Framework for C++

Version 1.2.x

s11n-devel@lists.sourceforge.net - <http://s11n.net>

ACHTUNG #1: this is a "live" document covering an in-development software library. Ergo... it may very well contain some misleading or blatantly incorrect information! Please help us improve the documentation by submitting your suggestions to our mailing list!

Abstract:

This document describes s11n (and "s11n-lite"), an object serialization framework for C++, version 1.2.x. It serves as a supplement to the s11n API documentation and source code, and is not a standalone treatment of the entire s11n library. Much of this documentation can be considered "required reading" for those wanting to understand s11n's features, especially its advanced ones.

s11n-lite, introduced in s11n version 0.7.0, simplifies the s11n interface, providing the features that "most clients need" for saving and loading arbitrary objects. It also provides a reference implementation for implementing similar client-side interfaces. The author will go so far as to suggest, with uncharacteristic non-humbleness, that s11n-lite's interface ushers in the *easiest-to-use, least client-intrusive, most flexible* general-purpose object serialization library ever created for C++.

Users who wish to understand s11n are strongly encouraged to learn s11n-lite before looking into the rest of the library, as they will then be in a good position to understand the underlying architecture and framework, which is significantly more abstract and detailed than s11n-lite lets on. Users who think they know everything about serialization, class templates and classloaders are *still* encouraged to *give s11n-lite a try*: they might just find that it's just too easy to *not* use!

Maintainer: stephan@s11n.net (list: s11n-devel@lists.sourceforge.net)

Table of Contents

1	Preliminaries	5
1.1	License	5
1.2	Disclaimers	6
1.3	Feedback	6
1.4	Credits	7
2	Introduction	9
2.1	Scope of this document	9
2.2	s11n's Dream	9
2.3	Main features	11
2.4	Notable Caveats (IMPORTANT)	12
2.5	WTF is s11n-lite?	13
2.5.1	Repeated warning: learn s11n-lite first!	13
2.6	Getting and installing s11n	14
2.6.1	Building under GNU systems	14
2.6.2	Building under Windows and "other" environments	14
2.6.3	Compiling and linking s11n client applications	14
2.6.4	Building under Cygwin, Mac OS/X (Darwin), etc.	15
2.7	Version Compatibility	15
2.8	Optional supplemental libraries	15
3	Main differences between 1.0.x and 1.2.x	16
3.1	s11n mantra change	16
3.2	Code consolidation and removal	16
3.3	Factory code reimplemented	17
3.4	node_traits<> changes, s11n::data_node replaced with s11n::s11n_node	17

3.5	New header conventions, faster compile times	17
3.6	Fetching class names of Serializables	18
3.7	Client-extendible s11nlite	19
3.8	~/s11nlite config file removed	19
3.9	Exceptions conventions	19
4	Core concepts	19
4.1	Terms and Definitions	19
4.2	The Official Grossly Oversimplified Overview of the s11n architecture	22
4.3	Process Overview	24
4.3.1	Serialization	24
4.3.2	Deserialization	24
4.4	Node Names and Property Key naming conventions (IMPORTANT!)	24
4.5	Overview of things to understand about s11n	25
4.6	Notes on error/success values (i.e., justifying the bool)	25
4.7	s11n and Patterns	26
4.7.1	The core	26
4.7.2	Classloader	27
4.7.3	Proxies	27
4.7.4	i/o	27
4.7.5	s11nlite	27
5	Serializable Interfaces: overview and conventions	27
5.1	Serialize Operator conventions	28
5.2	Deserialize Operator conventions	28
5.3	Data Node class names (IMPORTANT!)	28
5.3.1	Example of setting a node's class name	29
5.3.2	Using local library support for class_name()	29
5.4	Cooperating with other Serializable interfaces	30
5.5	Member template functions as serialization operators	30
6	Type Traits	31
6.1	s11n::node_traits<NodeType>	31
6.2	s11n::s11n_traits<SerializableType>	31
6.2.1	cleanup_functor	32
6.3	type_traits<T>	33
7	Five-minute intro: PODs and STL containers	33
7.1	#include ...	33
7.2	Saving	34
7.3	Loading	34
7.4	Now the really easy way: micro_api<>	34
8	How to turn JoeAverageClass into a Serializable...	35
8.1	Create a Serializable class	35
8.2	Specifying custom Serializable interfaces for InterfaceTypes	36
8.3	Specifying Serializer Proxy functors	36
9	How to turn JoeNonAverageClass into a Serializable...	37
9.1	JoeAverageClass<> class template	38
9.1.1	A cleanup functor	38
10	Doing things with Serializables	39
10.1	Setting "simple" properties	39
10.2	Getting property values	40
10.2.1	Simple property error checking	40
10.2.2	Saving custom Streamable Types	40
10.3	Finding or adding child nodes to a node	41
10.4	Serializing Streamable Containers	41
10.4.1	Trick: "casting" list or map types	41
10.5	De/serializing Serializable objects	42
10.5.1	Individual Serializable objects	42
10.5.2	Containers of Serializables	43
10.5.3	"Brute force" deserialization	43
11	Walk-throughs: implementing Serializable classes	44
11.1	Sample #1: Read this before trying to code a Serializable!	44
11.1.1	The data	44
11.1.2	The #includes	44
11.1.3	The serialize operator	44
11.1.4	The deserialize operator	45
11.1.5	Serializable/proxy registration	45
11.1.6	Done! Your object is now a Serializable Type!	45
11.2	Gary's code	46

11.2.1	Gary's Revelation	46
12	s11n registration & "supermacros" (IMPORTANT)	49
12.1	"Supermacros"	49
12.2	General: Interface Types	50
12.3	Choosing class names when registering	50
12.4	Registering Interface Types supporting serialization operators	51
12.5	Registering types which implement a custom Serializable interface	51
12.6	Registering Serializable Proxies	51
12.7	Where to invoke registration (IMPORTANT)	52
12.7.1	Hand-implementing the macro code (IMPORTANT)	53
13	Proxies, functors and algorithms	53
13.1	Commonly-used Proxies	53
13.1.1	I/OStreamable types: s11n::streamable_type_serialization_proxy	53
13.1.2	Arbitrary list/vector types: s11n::list::list_serializable_proxy	54
13.1.3	Streamable maps: s11n::map::streamable_map_serializable_proxy	54
13.1.4	Arbitrary maps: s11n::map_serializable_proxy	54
13.1.5	Arbitrary pairs: s11n::map::pair_serializable_proxy	54
13.2	Commonly-used algorithms, functors and helpers	54
13.3	When proxies aren't desired	55
13.4	Functor tags	56
14	Data Formats (Serializers)	56
14.1	General conventions	56
14.1.1	File extensions	57
14.1.2	Indentation	57
14.1.3	Entity translation	57
14.1.4	Magic Cookies	58
14.2	Overview of available Serializers	58
14.2.1	compact (aka, 51191011)	58
14.2.2	expatxml	59
14.2.3	funtxt (aka, SerialTree 1)	59
14.2.4	funxml (aka, SerialTree XML)	60
14.2.5	parens	60
14.2.6	sqlite3	61
14.2.7	simplexml	62
14.2.8	wesnoth	62
14.3	Tricks	63
14.3.1	Using a specific Serializer	63
14.3.2	Selecting a Serializer class in s11nlite	63
14.3.3	Multiplexing Serializers	63
14.4	Internals: flex's role in s11n	63
15	class_name() and friends	64
15.1	node_traits<T>::class_name()	64
15.2	s11n_traits<T>::class_name(const T *)	65
15.3	Class name of "unknown"	66
16	Exceptions conventions	66
16.1	The library throws when...	66
16.2	Throwing from client-side de/ser operations	67
16.3	Errors and SerT * deserialize<NodeT,SerT>(const NodeT &)	67
16.4	Exceptions and "external modules"	68
16.5	Specific guarantees	68
16.6	Making your Serializables exception-safe	69
17	SAM: Serialization API Marshaling layer	70
17.1	The SAM layer & interface	70
17.2	SAM's place in the API calling chain (and other important notes)	71
17.2.1	More about SAM<X*>	72
17.3	Historical changes	72
18	s11nlite specifics	72
18.1	Why use s11nlite?	72
18.2	client_api<NodeType>	73
18.3	File formats	73
18.4	Simple config files	74
18.5	micro_api<SerializableType>	74
19	Memory management and object relationships	75
19.1	Data nodes	75
19.2	Containers of pointers	75
19.3	Cleaning up before deserialization	76

19.4	Cleaning up after failed deserialization	76
19.4.1	Understanding the problem	76
19.4.2	Accommodating the problem, approach 1 (don't do this!)	77
19.4.3	Accommodating the problem, approach 2 (do this instead!)	77
19.5	Understanding "serialization ownership"	78
19.5.1	The basic case: objects own their own resources	78
19.5.2	Serializing pointers to data we don't own	78
19.5.3	Two-way parent/child relationships	79
19.6	Known problematic cases	80
19.6.1	std::set<T*> and std::multiset<T*>	80
19.6.2	std::map<K,V> and std::multimap<K,V>	80
20	Using plugins	80
20.1	Building plugins support	80
20.2	Win32 Achtung	80
20.3	The API	80
20.4	Basic Usage	81
21	s11n-related utilities	82
21.1	s11nconvert	82
21.2	s11nbrowser	82
22	Miscellaneous features and tricks	82
22.1	Saving non-Serializable	83
22.2	Saving application-wide state and Singletons	83
22.3	Saving lib state plus arbitrary client-specified state	84
22.4	"Casting" Serializables with s11n_cast()	85
22.5	Cloning Serializables	85
22.6	Half-intrusive proxying and useless friends	86
22.7	zlib & bz2lib support	86
22.8	Using multiple data formats (Serializers)	87
22.9	Sharing Serializable data via the system clipboard	87
22.10	Containers of const objects	87
22.11	Versioning of s11n data	88
22.12	Splitting up your output	88
22.13	Improving compile times	89
22.14	Know when you don't need to register a type to serialize it	89
22.14.1	Containers of Streamable types	89
22.14.2	Algos which don't need the s11n core API	90
23	Miscellaneous caveats, gotchas, and some things worth knowing	90
23.1	Serializing class templates	90
23.2	Cycles and graphs	90
23.3	Thread Safety	91
23.4	Polymorphic types and template parameters	91
23.5	Absolute No-nos (Worst Practices) for s11n[lite] client code	92
23.5.1	Do not change the name of a passed-in data node!	92
23.5.2	Do not use a single Data Node for multiple purposes!	92
23.5.3	Do not re-assign a reference returned by s11n::create_child()!	93
23.5.4	Do not use Serializers to implement classical iostream operator functionality!	93
23.5.5	Do not register a type as its own proxy!	94
23.6	Best Practices	94
23.6.1	std::auto_ptr<T> and s11n::cleanup_ptr<T>	94
23.6.2	Know when to use s11nlite and when not to	94
24	Functional serialization	94
24.1	#include	95
24.2	Example: serialize via std::for_each()	95
24.3	Composing custom algorithms from functors	96
24.4	Non-default-constructed proxies	96
25	Understanding the costs of deploying s11n	97
25.1	Learning curve	97
25.2	Intrusivity (or not)	98
25.3	Compilation costs	98
25.4	Memory/RAM costs	99
25.5	Runtime speed: s11n and the "Big O Notation"	100
25.6	Code maintenance costs	101
25.7	Money	101
26	Common problems	102
26.1	Satan speaks through the console during compilation	102
26.2	Containers serialize, but fail to deserialize	103

26.3	Abstract Interface Types for Serializables	103
26.4	Statically linking against s11n	103
27	Evangelism	104
27.1	Pointer/reference transparency for Serializables in the core API	104
27.2	Container-based algos which are pointer/reference-neutral	104
27.3	"Casting" between "similar" types	106
28	Comparing s11n and Boost::serialization	106
28.1	Cans and cannots	107
28.2	Compiler and platform portability	107
28.3	Archives vs S11nNodes	108
28.4	Non-intrusivity	108
28.5	Serialization of pointers	108
28.6	Data Versioning	109
28.7	API ease of use	110
28.8	Serialization Traits	111
28.9	Efficiency	111
28.10	The interesting part is...	112
28.11	In closing: s11n.net and Boost.org	112
29	Source tree innards	114
29.1	Build tree structure	114
29.2	Header file weirdness	114
29.3	Generated files	115
29.4	Plugins	115
30	In Hindsight...	116
30.1	The name "Data Node"	116
30.2	Patterns, formality, etc.	116
30.3	Exceptions	116
30.4	Build tree and code layout consistency	117
31	Is this the end?	117
32	Bibliography	118
33	Index	118

1 Preliminaries

1.1 License

"You cannot guaranty freedom of speech and enforce copyright law."

Ian Clarke

"This [document] is encrypted with ROT26 encoding. Decoding it is in violation of the Digital Millennium Copyright Act."

Anonymous Software Developer

The library described herein, and this documentation, are released into the Public Domain. Some exceptional library code may fall under other licenses such as BSD or MIT-style, as described in the README file and their source files.

All source code in this project has been custom-implemented, in which case it is Public Domain, or uses sources/classes/libraries which fall under LGPL, BSD, or other relatively non-restrictive licenses. It contains no GPL code, despite its "logical inheritance" from the GPL'd libFunUtil. Source files which do not fall into the Public Domain are prominently marked as such, and in absolutely no cases does this project use licenses which modify the license of code linked against it.

To be perfectly honest, i prefer, instead of *Public Domain*, the phrase *Do As You Damned Well Please*. That's exactly how i feel about sharing source code.

Whatever the license, however, i will request that if you redistribute your own libraries based off of this code, please do not use the same *installed* binary/library/header file names. For example, if you redistribute libs11n, please do not install the library as libs11n.so, nor the headers under `<s11n.net/s11n/...>`. Doing so will inherently complicate cases where both of our copies of s11n are used on the same systems.

1.2 Disclaimers

"This information provided free of charge for those willing to accept it. Others who wish to be spoon-fed may acquire my services at the discounted rate of 235 Euro per hour or part thereof."

Anonymous Software Developer

The obligatory disclaimers include:

1. This manual will make *no sense whatsoever* to most people. It is target at experienced C++ programmers ("intermediate level" and higher), and makes many assumptions about prior C++ knowledge.
2. *Don't let the size of this manual make you think that using s11n is difficult!* Using s11n (*especially s11n-lite*) is simple and straightforward, even for non-guru C++ coders. It also has a number of "power user" features which can be exploited by those who truly understand the architecture.
3. There is admittedly a lot of hype and evangelism in this manual, but i personally believe it to all be justified.
4. s11n is continually under development and is constantly being tweaked. The basic model it is based on has proven to be inordinately effective and low-maintenance since it was introduced in the QUB project (qub.sourceforge.net) by Rusty "Bozo" Ballinger in the summer of 2000. This implementation refines that model, vastly expanding its capabilities.
5. This software and documentation are PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE.
6. Reading disclaimers makes you go blind. ;)
7. Writing them is even worse. :/
8. This list of disclaimers might not contain all the necessary disclaimers.

And, finally:

This library is developed in my private time and the domain and web site are funded by myself. With that in mind: unless i am kept employed, this project may "blink out" at any time. That said, this particular project holds a special place in my heart (obviously, or you wouldn't be seeing this manual and all this code), so it often does get a somewhat higher priority than, e.g. dinner or lunch. Should you feel compelled to contribute financially to this project, please do so via the donation program hosted by SourceForge and PayPal:

https://sourceforge.net/donate/index.php?group_id=104450

Donations will go toward keeping the web site online and the domain name registered, and potentially to cover internet access fees. If anyone is interested in providing a grant to this project, please contact us directly. We would be thrilled. Of course, non-financial contributions, e.g. code, documentation, and bug reports, are of course also welcomed.

1.3 Feedback

"Like a Kenny Loggins record, no one's ever gonna hear ya."

Bloodhound Gang

The s11n project's home page is:

<http://s11n.net/>

The author is stephan beal (stephan@s11n.net). Feel free to contact me directly, but i would ask that questions about the library be directed to our development mailing list:

s11n-devel@lists.sourceforge.net

You do not need to subscribe to the list in order to post there.

By all means, please feel free to submit feedback on this manual and the library: positive, negative, whatever... as long as it's constructive it is always happily received. While few who know me would say that i am a pedantic person, i am extremely pedantic when it comes to documenting software: if you find any errors or gaping holes in these docs, please point them out!

If this gives you any idea of how seriously feedback is taken:

- The whole 0.7.0 rewrite, and the abstractions and simplifications which grew out of it, were triggered by

Ton Oguara's feedback about his problems serializing class templates. That is indeed a deceptively tricky problem, and the older code could only handle non-trivial cases with a non-trivial amount of code generation. The 0.7 framework can do this with "relative" ease, and 0.8+ makes it trivial in many cases.

- This particular document (the one you're reading now), was largely inspired by **Gary Boone's** feedback on the difficulties of getting started with s11n. Also, the changes in the registration processes from 0.7x to 0.8 were inspired by Gary.
- s11n-lite was developed largely because of Ton's and Gary's feedback.
- The massive build tree re-orgs between 0.8.x and 0.9.x were inspired by the Debian Project's **martin f. krafft** (yes, he prefers it spelled lower-case).

The contact address, should you also feel compelled to write what you *really* think about s11n, is at the top of this document.

Now, i can't promise to rewrite everything every time someone wants a change, but all input is certainly considered. :)

Whatever it is you're trying to save, s11n *wants* to help you save it, and goes through great pains to do some deceptively difficult tricks to simplify this process as much as practically possible. If it *can't* do so for your cases, then please consider helping us change s11n to make it capable of doing what you'd like it to. It is my firm belief that the core s11n framework can, with very little modification, save *anything*. What is currently missing are the algorithms which may further simplify the whole process, but only usage and experimentation will reveal what that toolkit needs to look like. If you come across some great ideas, please share them with us!

:)

1.4 Credits

"It's a thankless job, but I've got a lot of Karma to burn off."

Anonymous Software Developer

There is no single, complete list of all people who have influenced this project. A partial list, in no particular order:

(If i have left *you* off of the list, please let me know!)

- My mother and step-father **Bonnie & David Pickartz**, my father and step-mother **Joseph & Gail Hudgins**, my step-mother-by-adoption **Kathy Beal**, and my belated adopted dad, **Gerry Beal**¹. They just all need to be thanked in general. The 0.9.x branch, from 0.9.4 to 1.0.0, was directly funded by a *very gracious* donation from the Pickartz family.
- **Rusty "Bozo" Ballinger** wrote the conceptual forefather of s11n (<http://libfunutil.sourceforge.net>). (Rusty, if you're still out there, *get in touch!*) As far as i know, Rusty also coined the phrase "s11n" as a short form of "serialization", which i then stole as a domain name.
- **Ton Oguara** accidentally inspired the whole 0.6 -> 0.7 rewrite/refactor by showing me how much client-side effort was *really* needed to de/serialize class templates.
- **Gary Boone** provided valuable feedback on a range of documentation and features, particularly on making it easier for developers to get started with s11n. Many of the 0.8.x improvements exist because of Gary's feedback. Gary also is credited with coming up with a useful naming convention for Serialization Proxies, `MyType_s11n` (a convention now adopted in many of my projects).
- **Roger Leigh** provided the information needed to add `libltdl` support to the classloader.
- **Tom**, from `comp.lang.c++.`, provided an interesting fix for an "interface annoyance" in the classloader. It is still used to this day in registering class factories.
- **martin f. krafft**, of the Debian Project, put in a great deal of effort to get the 0.8.x series into the Debian, and was the driving force behind the 0.8.x -> 0.9.x source tree re-orgs. His continued feedback is always insightful.
- **Marshall Cline**, of C++ FAQ fame, helped to correct some of the errors in the documentation regarding cycles, joins and trees. His FAQ has a great section on the topic of serialization in C++: <http://www.parashift.com/c++-faq-lite/>
- **Robert Ramey**, author of the Boost serialization library (<http://www.rrsd.com>), for several insightful email conversations on the topic of serialization. His well-crafted library is compared to this one (we might even say praised) in some detail in section 28.
- **Steve Madere** suggested adding unit tests to the source tree, and within an hour of doing so 2

¹ i've got 8 brothers and 4 sisters. Yes, i actually *do* know all of their names: (in no particular order) Toby, Gerald, Ty, Trevor, Teven, Wayne, Wesley, David, Margorie, Melisa, Ashley and Cindy (though i've never actually met Cindy). Their birthdays? Err.... ???

significant bugs were caught and fixed. He also made a financial contribution via the SourceForge/PayPal donation system.

- **Andreas Jochens** provided several patches for compiling the 0.8.x tree under gcc 3.4.
- **Mike Radford** provided more patches for gcc 3.4 and gave me ssh to his box to let me fix a couple more.
- **Patrick Lin** demonstrated and helped localize a long-standing container-of-containers deserialization bug-in-waiting which couldn't wait any longer on his machine (existed from 0.8.x until 0.9.17).
- **Keven Weber** helped track down a couple bugs by allowing me ssh access into his machine, where the bugs were appearing.
- **Christian Prochnow**, project lead of P::Classes (<http://pclclasses.com>), allows me to integrate s11n support into P::Classes 2.x. The port provides a great opportunity for bug finding and cleanups.
- **Dr. Marc Duerner**, first for inviting me to pclclasses.com and secondly for his continued and on-going feedback and hacking sessions.
- **Gregor Jehle**, also of pclclasses.com, reported compile problems on AMD64 and allowed me ssh access to his box to track down and fix them.
- **Ashran**, my personal best friend and author of the *Hackersquest* Everquest(tm) emulator (<http://hackersquest.org>), was the first to compile s11n under Windows.
- **Peter "What's Happ'nin'!!?!?!?" Angerani**, my long-time friend and mentor, for his continued support and feedback.
- To my esteemed Unix-loving colleagues, **Ralf Lehmann** and **Martin Tessun**, for agreeing, after bribing them with their own personal Easter Egg in this document, to look over this manual for me. (*Hier ist euer Easter Egg, Jungs!!!*)
- **SourceForge** (<http://sourceforge.net>) has been hosting my code since 2000, and without them s11n would have neither mailing lists, a bug tracking site, nor a public CVS tree. i encourage all users of SourceForge to support their service by buying a yearly subscription to their site.
- **Pete Harlow** pulled out a can of OpenDocument wizardry and ported the 120-page 1.2.0 library manual from Lyx format to OpenDocument.

Various published authors have, rather unknowingly, had *profound* impacts on various design decisions during s11n's evolution:

- **Scott Meyers** - a *huge* percentage of my code is influenced by Scott's always-practical advice. All of his books must be on any C++ coder's bookshelf. *Here's your biggest fan, Scott!*
- **Andrei Alexandrescu** - his *Modern C++ Design* was the necessary catalyst i needed for realizing the classloader implementation, and provided the basis for the internals of the `phoenix<>` class, which is used *extensively* by s11n.
- **Herb Sutter** - A couple of his (very numerous) articles have led to direct changes in this library. e.g. a breaking-down of some of the member-based interfaces into free functions was inspired by his "What's in a class?" article.
- **Stephen Dewhurst**, author of C++ *Gotchas*: every time i write "template class" and correct it to "class template", or change the word "method" to "function", i think of Stephen. ;) If i recall correctly, Stephen also introduced me to the idea of Monostates, which are conceptually similar to what i've been calling "Context Singletons."
- C++ *Templates: The Complete Guide*, by **Nicolai M. Josuttis** and **David Vandevoorde**, as well as Josuttis' *The C++ Standard Library*, were instrumental in implementing much of the template code used by this library. The latter is always the first book i reach for when i've got a question about the STL, and 99% of the time it has the answer².

i try to keep keep the list of contributors up-to-date via an RSS feed:

<http://s11n.net/rss/s11n-contributors.xml>

2 Introduction

So you want to save some objects? Strings and PODs³? Arbitrary objects you've written? A `FooObject` or `std::map<int, std::string>` or `std::list<MyType*>`?

What?!?!? You've got a:

```
std::map< int, std::list< std::map< double, FooObject<X *> * > > >
```

?!?!?

² i say 99% because i generally mistrust statements which include a "100%" qualifier, but the truth is i can't remember a time when this book didn't have what i was looking for.

³ Plain Old Data types, such as `int`, `char`, `bool`, `double`, etc.

s11n is here to Save Our Data, man!

Historically speaking, saving and loading data structures, even relatively simple ones, is a deceptively thorny problem in a language like C++, and many coders have spent a great deal of time writing code to serialize and deserialize (i.e., save and load) their data. The s11n framework aims (rather ambitiously) to completely end those days of drudgery.

s11n, a short form of the word "serialization"⁵, is a library for serializing... well, just about any data structure which can be coded up in C++. It uses modern C++ techniques, unavailable only a few years ago, to provide a *flexible, fairly non-intrusive, maintenance-light, and modern* serialization framework... *for a programming language which sorely needs one!* s11n is particularly well-suited to projects where data is structured as hierarchies or containers of objects and/or PODs, and provides *unprecedentedly simple* save/load features for most STL-style containers, pretty much regardless of their stored types.

In practice, s11n has far exceeded its original expectations, requirements and goals, and it is hoped that more and more C++ users can find relief from Serialization Hell right at home in C++... via s11n.

A brief history of the project and a description of its main goals are available at:

<http://s11n.net/history.php>

2.1 Scope of this document

Originally, this document set out to provide a quick-start guide to using the library's main features. Over time it has evolved to cover nearly every aspect of the library. Between this manual, the API documentation, and the sample code provided with the library, pretty much all of your questions about the library should be answered. If not, feel free to email us with your questions.

As always, the sources are the definitive place for information. That said, i'm a firm believer that developers should not have to read the sources in order to be able to use a library, so there is an absurd amount of documentation.

2.2 s11n's Dream

Anyone who has had to hand-code save and load support for their data, even if only for relatively trivial containers and data types (e.g. even non-trivial strings), will almost certainly agree with the following statement:

Saving data is relatively easy. Loading data, especially via a generic interface, is *mind-numbingly, ass-kickingly difficult!*

The technical challenges involved in loading even relatively trivial data, *especially* trying to do so in a unified, generic manner, are *downright frigging scary*. Some people get their doctorates trying to solve this type of problem⁶. Complete *branches* of computer science, and hoards of computer scientists, students, and acolytes alike, have researched these types of problems for practically *aeons*. Indeed, their efforts have provided us a number of *critical* components to aid us on our way in finding the Holy Grail of serialization in C++...

In the 1980's IOSTREAMS, the predecessor of the current STL iostreams architecture, brought us, the C/C++ development community, *tremendous* steps forward, compared to the days of reading data using classical brute-force techniques, such as those provided by standard C libraries⁷. That model has evolved further and further, and is now an instrumental part of almost any C++ code⁸. However, the practice of directly manipulating data via streams is showing its age. Such an approach is, more often than not, not suitable for use with the common higher-level abstractions developers have come to work with over the past decade (for example, what does it *really* mean, semantically speaking, to send a UI widget to an output stream?).

In the mid-1990's HTML became a world-wide-wonder, and XML, a more general variant from same family of

4 The only [remaining] inherently difficult part for this one is getting the proper type *names* for each component of the container hierarchy! This problem discussed at length in this documentation, the s11n sources, and the class_loader library manual. It's not as straightforward as it may seem. Interestingly, for many cases (non-polymorphic types) we can actually get by without knowing the type's name.

5 's11n' was coined by Rusty Ballinger in mid-2003, as far as i am aware. It follows the tradition set by "i18n", which is short for "internationalization" - the number represents the number of letters removed from the middle of the word.

6 But all i got was this library manual. ;)

7 That was all well before my time, but i read a lot of C++ books. ;)

8 Are you going to tell me you never use std::cout and std::cerr? Yeah, right. Tell it to your grandma - maybe she'll believe you.

meta-languages HTML evolved from, SGML⁹, leapt into the limelight. Practically overnight, XML evolved into *the* generic platform for data exchange and, perhaps even more significantly, *data conversion*. XML is here to stay, and i'm a *tremendous* fan of XML, but XML's era has left an even more important legacy than the elegance of XML itself:

More abstractly, and more fundamentally, the popularity and "well-understoodness" of XML has *greatly* heightened our collective understanding of abstract data structures, e.g. DOMs [Document Object Models], and our understanding of the general needs of data serialization frameworks. *These points should be neither overlooked nor underestimated!*

What time is it now? 2004 already? It looks like we're ready for another 10-year cycle to begin...

We're in the 21st century now. In languages like Java(tm) and C# serialization operations are basically built-in¹⁰. Generic classloading, as well, is EASY in those languages. Far, far away from Javaland, the problem domain of loading and saving data has terrified C++ developers *for a full generation!*

s11n aims, rather ambitiously, to put an end to that. The whole general problem of serialization is a very interesting problem to me, on a personal level. It fascinates me, and s11n's design is a direct result of the energy i have put into trying to rid the C++ world of this problem *for good*.

Well, okay, i didn't honestly do it to save the world[']s data]:

i want to save my objects!

That's my dream...

Oh, my - what a coincidence, indeed...

That's s11n's dream, too...

s11n *actively* explores *viable, in-language* C++ routes to *find*, then *take*, the C++ community's *next major evolutionary step* in general-purpose object serialization... all right at home in ISO-standard C++. This project takes the learnings of XML, DOMs, streams, functors, class templates (and specializations), Meyers, Alexandrescu, Strousup, Sutter, Dewhurst, PHP, "Gamma, et al", comp.lang.c++, application frameworks, Java¹¹, and... even lowly ol' me (yeah, i'm the poor bastard who's been pursuing this problem for 3+ years ;), and attempts to create a unified, generic framework for saving... well, damned near anything. Actually, *saving* data is the easy part, so we've gone ahead and thrown in *loading* support as an added bonus ;).

In short, s11n is attempting to apply the learning of an entire generation of software developers and architects, building upon of the streets they carved for us... through the silicon... armed only with their bare text editors and the source code for their C compilers. These guys have my *utmost* respect. Yeah, okay... even the ones who chose to use (or implement!) `vi`. ;)

Though s11n is quite young, it has a years-long "conceptual history"¹², and its capabilities *far, far* exceed any original plans i had for it. Truth be told, i use it in *all* of my C++ code. i can finally... *finally*, **FINALLY SAVE MY OBJECTS!!!!**

i hope you will now join me in screaming, in the loudest possible volume:

It's about damned time!!!

2.3 Main features

"I don't make my mistakes more than once. I store them carefully and after some time I take them out again, add some new features and reuse them."

Anonymous Software Developer

For the most part, the features list is the same as for s11n 1.0.x. For those of you who haven't used 1.0.x, the library's primary features and points-of-interest are:

- Quite possibly the *most flexible* and *easiest-to-use* C++ serialization framework *in the known*

⁹ [Standard,Structured] Generic Markup Language.

¹⁰ Though i do have very deep fundamental differences with Java's built-in serialization model!

¹¹ Incidentally, not C#: s11n was started before i ever touched C#. In all honesty, i find C#'s core model to be inferior to s11n, at least in terms of its client-side interface. For example, it really bugs me that in C# (or any other serialization framework), the client must know something so basic as what file format their data is stored in. i say (and s11n says): only a file's i/o parsers *really* care what format a file is in

¹² Utility-class coding, and *lots* of design thought, started in early 2001. The "real coding" began in September, 2003, once i finally cracked the secrets i needed to implement the classloader.

universe.¹³

- Provides client code with easy de/serialization of arbitrary streamable types and user-defined Serializable types.
- Out of the box it supports all standard STL containers: `std::list`, `vector`, `set`, `multiset`, `map`, `multimap` and `valarray`.¹⁴
- Lends itself well to a large number of uses, from de/serializing arbitrary vectors or maps of data (a-la config files) to saving whole applications in one go.
- Does not tie clients to a specific Serializable interface/hierarchy. The internally-used interfaces can be *easily* directed to use client-specific interfaces, which need not even be virtual. This means that the library's interface can be made to conform to client-side objects' needs, as opposed to the other way around.
- Serializable Proxying allows clients to attach proxy classes to arbitrary types, such that the proxy type is delegated all de/serialization operations. The end result is that it is possible to serialize a given type without having to touch a line of that type's code, nor does that type have to know its playing along.
- Advanced techniques allow client code to completely reimplement/replace most of the library's underlying layers with their own - without touching the s11n code. For example, class factories or even the client-to-core API translation layer can be replaced by providing certain class template specializations.
- Integration into existing class hierarchies is straightforward, quick, relatively painless, and can often be incrementally applied to subsets of a project over time, as needed, as opposed to forcing a client to completely refactor. In fact, using proxies means client classes don't normally have to change *at all* to be transformed into "True Serializables."
- The data persistence model inherently does not suffer (as, e.g. Java's does) from the problem of invalidating serialized data every time an internal change is made to a Serializable data type. It's "structure-and-properties"-based system ensures that legacy data do not become invalid until developers¹⁵ *want* them to become so.
- It sports *compile-time type-safe classloading* without the use of a *single* type-cast (neither in the client nor in the library). The classloader is factory-based, and can load just about any classes, including 3rd-party classes, without them knowing they are participating. Transparently loading new types from DLLs is supported if available on your platform.
- The API is *100% data-format agnostic* and places no file naming conventions client data files. Several (err... many) different data format handlers currently exist, and adding custom Serializers is fairly painless: all you need is an input parser and an output formatter¹⁶. Existing formats include three XML dialects, one MySQL-powered "format", and experimental add-on support for ftp/http which works with arbitrary file-based formats. That is, as far as i am aware, more formats than any existing serialization library, regardless of implementation language. Why so many? Mainly to show that it can be done ;).
- Does not impose any special file name conventions or restrictions on clients¹⁷. That is, if you want to call your saved data `MyData.doc`, go right ahead.
- *All* clients of s11n may share serialized data between themselves, regardless of their underlying client serialization interfaces. If their APIs can see each others' factories then they can also transparently fully deserialize each others' data.
- Optional client-transparent zlib and bz2lib file de/compression, for 60-95% file size reduction. When enabled, de/compression happens transparently - usage of s11n does not change one iota.
- The i/o sub-framework is stream-centric, not file-centric. This sub-module is effectively optional: clients are not required to use *any* of the supplied i/o code, but must then supply their own Serializers (i/o handlers, which need not use streams, but could use a relational database or any other back-end).
- The primary data structures follow STL [Standard Template Library] conventions and are container/functor/algorithm-centric, thus many generic algorithms can be easily applied to them. The library comes with several useful functors and algorithms for working with serialized data. This also allows complete separation between the processes of the state storing/restoration and any resulting i/o.
- Uses only ISO-standard C++ constructs, no compiler-specific extensions.
- Allows clients *complete* control over *how* an object is serialized: s11n makes no assumptions about

13 On a features/technical level, the only currently-existing C++ serialization framework which can even begin to compare with s11n is Dr. Robert Ramey's Boost serialization lib, available via <http://www.boost.org>. For a comparison of Boost and this library, see section 28.

14 Reminder: `std::queue`, `deque` and `stack` are not strictly containers - they are *container adapters*. The unusual traversal requirements of queues and stacks make them difficult to serialize efficiently.

15 Or, admittedly, the all-powerful Marketing Director.

16 A new Serializer can be implemented in under an hour if one has related Serializer or parser code to start from, and can normally be done in as little as a few hours even when writing from scratch. The real effort is normally in writing the input parser: the only special consideration normally needed is the escaping of, e.g. strings (this is format-dependent).

17 It might be limited by your underlying filesystem or STL, e.g. in regards to Unicode. s11n has no special support for Unicode, relying on `std::string` for all string operations.

what you want, it only tries (very hard) to help you meet your data persistence needs. That said, s11n can be told how to serialize many complex object types with very little instruction, so clients need not normally do very much work.

- It comes with an *absurd* amount of documentation, in the form of this document, the API docs, and the web site.

Okay, okay, we'll stop there! ;)

2.4 Notable Caveats (IMPORTANT)

It would be dishonest (even if only mildly so ;) to say that s11n is a magic bullet - the solution to all object serialization needs. Below is a list of currently-known major caveats which must be understood by potential users, as these are type types of caveats which may prove to be deal-breakers for potential s11n users. Much more detailed information and speculation about the overall client-side costs of deploying s11n-based code can be found in section 25.

- As it is heavily based on class templates, it is implemented largely as inlined code in header files (for complex linking reasons). The end effect on clients is that compilation times and object/binary file sizes *do* suffer. (One user reports that compile times increase by as much as 14 *times* when building with libs11n 0.8.x, but this has been cut drastically since his report.) Some code is in implementation files, so clients must still link to the s11n library, just as they would for any typical C/C++ library.
- Due largely to the side-effects of heavy reliance on class templates, s11n is unsuitable for systems with very limited filesystem space or main memory (e.g. embedded systems, handheld computers, etc.).
- s11n, at its core, can be quite difficult to grasp. It's not the details which are difficult for most people, i think, but the fact that the details are hidden behind very abstract "conventions" and "close approximations". Using the s11nlite interface will completely eliminate most potential "startup problems" when getting used to this library. What is s11nlite? See section 2.5.
- s11n can serialize, *but not deserialize*, classes containing references. There are workarounds, but they require modifying such classes to internally hold a pointer instead of a reference, making them default constructable, and maybe other minor changes.
- The supplied build tree will only run on GNU-based systems. That is, systems running all the common GNU tools like `make`, GNU `bash`, and other exceedingly common Open Source tools, like `perl`. That said, the code itself should be easily portable to other build systems, so long as those hosts support appropriate compilers (see below). We will gladly host build-related files for other platforms or build environments (e.g. GNU Autotools, Microsoft environments, etc.) in the distribution and/or web site, should users submit those.
- Requires a relatively recent, ISO-conformant C++ compiler with excellent support for class templates. Only known to work with GCC 3.2x - 3.4.x, and known to *NOT* work with GCC 2.9x. On Win32 platforms, as of version 1.1.2 it is known to build under MSVC 2003 and 2005.
- s11n is untested with serializing binary data. It "should be possible", but implementing it in terms of the current Serializers (e.g. as string-encoding conversions like base64) would be rather inefficient, i think (even more so than s11n's normal techniques, i mean). That said, any data which can ultimately be represented as a one or more `std::string` objects and can be structured in a DOM-like fashion (even if only via transformation) should pose no problems at all for s11n. (We avoid binary formats partially so that we can evade the problems related to machine endianness.)
- The library currently has no algorithms for saving *graphs* - that is, *structures with joins*. This *can* and *has* been done in s11n, but no generic algorithms are (yet) provided for doing so. For more information see section 23.2.
- s11n is untested in multi-threaded environments. See section 23.3 for more details and speculation.
- It is driven with Generic Programming and reusability/maintainability in mind, not High-performance Computing, and thus it may not be performant enough for projects which need, really, *really* fast code. (That said, s11n is acceptably fast for all uses i've had for it. Try it out and make your own judgement.) Its general model inherently at-least-linear (or even worse), as discussed in more detail in section 25.5.
- s11n's development is primarily steered by my hobbies and my client-side needs, and is constantly under experimentation.
- s11n will not work when built as a static library, for complex linking purposes.

2.5 WTF is s11nlite?

(WTF is a technical term used very often by I.T. personnel of all types. It is short for "*What the foo?!?*")

s11nlite is a "light-weight" s11n sub-interface written on top of the s11n core and distributed with it. It provides "what most clients need for serialization" while hiding many of the details of the "raw" core library from the client (trust me - you *want* this!). Overall it is *significantly* simpler to use and, as it is 100% compatible with the core, it still has access to the full power "under the hood" if needed. s11nlite also offers a potential starting

point for clients wishing to implement their own serialization interfaces on top of the s11n core. Such an approach can free most of a project's code from direct dependencies s11n by hiding serialization behind an interface which is more suitable to the project. (Such extensions are beyond the scope of the document, but feel free to contact the development list if you're interested in such an option, and we'll help you out.)

Historically, the s11n architecture has been significantly refactored three times, and it has evolved to be more and more useful with each iteration. This particular iteration is light years ahead of its predecessors, in terms of power and flexibility, and is also much simpler to work with and extend than earlier architectures.

Users new to s11n are *strongly* encouraged to learn to use the code in the s11nlite namespace before looking into the rest of the library. Doing so will put the coder in a good position to understand the underlying s11n architecture later on. Users who think they know everything are still encouraged to give s11nlite a try: they might just find that it's just too easy to *not* use! Don't let the 'lite' in the name *s11nlite* fool you: it's only called s11nlite because it's a *subset* (but a functionally complete one) of an even more powerful, more abstracted layer known as "the s11n core" or "core s11n."

2.5.1 Repeated warning: *learn s11nlite first!*

We'll say this again because people don't seem to want to believe it...

i wrote s11nlite because i, the author of s11n, found s11n's core "too detailed" for client-side use. i like the general core model, but it is cumbersome to use directly, due to the many places where template parameter types must be specified. So i got tired of dealing with it and sought out a Simpler Way of Doing Things. That is what s11nlite is all about.

If you think i'm kidding about learning s11nlite first, take a look at this note from s11n user Paul Balomiri¹⁸:

"I didn't trust you on the point about understanding s11lite first (don't ask why, it was a mistake anyway)."

That is, for the *vast majority of cases*, s11nlite provides everything clients need as far as using s11n goes, and has a *notably* simpler interface than the core library. s11nlite, combined with the various generic serialization algorithms shipped with s11n (e.g. in `listish.hpp` and `mapish.hpp`), provide a complete interface into the framework.

Another point to consider: in client-side code i (s11n's author) generally use s11nlite and the generic algos/proxies, and rarely dip down into the core, nor do i deal with the Serializer interface from client code. Thus, i can assure you - a potential s11n client - that s11nlite can do almost anything you'd want to do with this library, and is *significantly easier* to work with than the core interface is.

If you still don't believe me, please re-read this section until you do.

2.6 Getting and installing s11n

"Linux sucks twice as fast and 10 times more reliably, and since you have the source, it's your fault."

Anonymous Software Developer

s11n can be downloaded from:

`http://s11n.net/download/`

2.6.1 Building under GNU systems

The build tree shipped with the main source tree is GNU-centric, because i happen to use GNU tools. Building it on systems which do not host GNU tools (gcc, make, bash, etc.) will require creating custom build control files (project files, makefiles, or whatever).

To build the library, use the conventional approach:

```
./configure [--options ...]
make
```

¹⁸ As of this writing, Paul uses s11n 1.0.x for some massive data sets: 10 million data points describing the whole street network of Vienna, Austria. :)


```
make install
```

The most common option passed to configure is `--prefix=/some/path`, which defines the top-level path for installing the library. If you do not have admin rights on the machine, i suggest using `--prefix=$HOME`, and adding `$HOME/lib` to your `LD_LIBRARY_PATH`.

Pass `--help` to configure for a list of more options.

2.6.2 Building under Windows and "other" environments

"People say it is hard to switch from Windows to UNIX; sure: but it is impossible to switch from UNIX to Windows!"

Anonymous Software Developer

Starting with release 1.2.0, we release a "static" variant of the source tree which comes with all generated files pre-generated and doesn't include any build-related files except for a very simple Makefile. The intention is to make it possible to easily pull the s11n source code into your own build tools, regardless of the platform. To download one of these releases, look for s11n releases named `libs11n-VERSION-nobuildtools.*`.

As of version 1.1.2, s11n is known to compile under at least a couple variants of MS Dev Studio. For full instructions on building under Windows see the file named `README.WIN32`, which comes with the source distribution. The demo Makefile will also be helpful, as it shows which sources belong to which parts of the library.

2.6.3 Compiling and linking s11n client applications

On Unix systems, use the `libs11n-config` script, installed under `PREFIX/bin`, to get information about your `libs11n` installation. This includes compiler and linker flags clients should use when building with s11n. It may (or may not) be interesting to know that `libs11n-config` is created by the configure process, so if you have used a build process other than the one shipped with the library, you may not have this script, or may need to generate it by hand.

When linking client binaries and shared libraries on Unix systems, you must use the `-rdynamic` (or equivalent) linker option. If you do not, factory registrations will not work (they will never happen) and deserialization of pointer types will therefore fail. This is unfortunate, but true.

As with all Unix binaries which link to dynamically-loaded libraries, clients of `libs11n` must be able to find the library. On most Unix-like systems this is accomplished by adding the directory containing the libs to the `LD_LIBRARY_PATH` environment variable. Alternately, many systems store these paths in the file `/etc/ld.so.conf` (but editing this requires root access). To see if your client binary can find `libs11n`, type the following from a console:

```
ldd /path/to/my/app
```

Example:

```
stephan@owl:~/cvs/s11n.net/1.1/s11n/src/client/sample> ldd ./demo_coord
linux-gate.so.1 => (0xfffffe000)
libs11n.so.1 => /home/stephan/cvs/s11n.net/1.1/s11n/src/libs11n.so.1
(0x40019000)
...
libdl.so.2 => /lib/libdl.so.2 (0x4034d000)
```

If you see a message like "not found" next to a library, then the dynamic linker cannot find it. In that either you do not have the library or it is not in one of the search paths used by your system's dynamic library loader, which are typically defined in the environment variable `$LD_LIBRARY_PATH` or the file `/etc/ld.so.conf`. If you modify `ld.so.conf`, you will need to run `ldconfig`, as the root user, for the changes to take effect.

2.6.4 Building under Cygwin, Mac OS/X (Darwin), etc.

As i not have these tools, i cannot directly do ports to them. Anyone interested in assisting, please get in touch.

The source code is believed to be compilable under any recent, standards-compliant C++ platform. It might require a tweak here and there for specific platforms, but no major incompatibilities are expected.

When porting to a non-Unixish, non-Win32 platform, you will either need to disable plugins support altogether, use the no-op (default) plugin handler, or add an implementation for your platform (see `src/plugin/plugin.*.cpp`).

2.7 Version Compatibility

*"In this library, the only thing which is constant is the namespace."
Anonymous Software Developer*

As of the release of 1.0.0, libs11n will attempt to follow the version compatibility guidelines laid out below.

- **Major version number:** the X in X.Y.Z. With Major version increments there are no set guidelines as to what might change, and there are absolutely no guarantees of compatibility with older releases.
- **Minor version number:** the Y in X.Y.Z. Minor number increments may or may not be API-compatible with previous releases. As per "the Linux convention", odd-numbered Minor numbers represent "development trees", intended for developers and early-adopters. Likewise, even-numbered Minor numbers represent "stable" trees, suitable for client use. Within development trees, existing conventions might be changed significantly at any time, whereas in stable trees they will not.
- **Patch level:** the Z in X.Y.Z. Patch-level changes should be conventions-compatible with earlier releases in the same Minor number, and preferably binary-compatible. Binary compatibility will be sacrificed in the interest of "important" fixes or additions, but this should be the exception, not the rule. Within the same *even* Minor number, well-established conventions will never be drastically altered from one patch level to the next (in development trees, anything goes).

s11n's basic model ensures that *data formats are almost always compatible across differing s11n versions*, and that when they are not then it was intended to be so (it doesn't happen by accident). It is very rare that a *format* ever changes after its initial definition, and thus data saved with s11n are "almost guaranteed" to be compatible across s11n versions, assuming a given format is not abandoned at some point. In cases where such compatibility is broken, i will do my best to release a tool to convert older data files to newer formats. Historically speaking, only once has an s11n-supported format ever changed significantly after its initial release (and two of them have stayed the same since the year 2000). See section 14.2 for more information on the available Serializers.

2.8 Optional supplemental libraries

s11n can make use of the following additional libraries, but does not strictly require them:

- `zfstream`, a published-by-s11n.net lib, provides transparent de/compression for files using zlib and bz2. This library comes as part of the source bundle but is not required by s11n 1.1 and higher (it is required in 1.0.x). Direct dependencies on this library are not recommended, as this library will be replaced once i get my hands on some more flexible code being written by my friend Marc Duerner. If you want zlib/bz2lib compression *now*, however, this is the way to plug it in to s11n.
- `libexpat`, required only if you want to build and use the expat-based XML Serializer (section 14.2.2). This library is almost certainly installed on almost all Unix-like OSes, because it is the *de facto* standard amongst the various Open Source, C-based XML libraries. The configure script checks for it, and disables the expat-based Serializer if the expat library and headers are not found.

The source trees shipped without build files will have these options turned off in the main configuration header(s).

3 Main differences between 1.0.x and 1.2.x

"We're going to tell people that even if (it) means we're going to break some of your apps, we're going to make these things more secure. You're just going to have to go back and fix it."

Craig Mundie, of Microsoft,

<http://www.wired.com/news/technology/0,1282,56381,00.html>

This section will only be of interest to users of s11n 1.0.x, and summarizes the significant changes from that version (i.e., those which would directly affect users of 1.0). This entire section assumes prior knowledge of how s11n works. If you have never used 1.0, and are just starting out with s11n, skip this section entirely - it is likely of no value to you unless you're a fan of arcane software history. New users are strongly recommended to go straight to 1.2.x, bypassing 1.0 altogether.

While this section might look quite large, architecturally *very little* has changed since 1.0. However, there have been a number of code reorgs and a few relatively low-impact additions. It is believed that porting from 1.0 will require relatively little client-side work (but some will be required, mainly due to header changes).

3.1 s11n mantra change

Since the beginning, s11n's core mantra has been that s11n is here to *Save Your Data, man!* As it turns out, that is a misrepresentation. Actually... it's a bald-faced lie. The honest truth is that s11n is here to...

Save Our Data, man!

Note the one-letter change, which is more significant than the single missing letter might imply.

3.2 Code consolidation and removal

One of the major goals of 1.1 is to have a tree which will compile on (Microsoft(tm) Windows(tm))(tm) platforms. Another is simplifying support for arbitrary build processes. Yet another related goal is to make the core library more easily forkable, so as to be able to copy it into arbitrary trees.

One requirement for achieving these is some major code refactoring, mainly elimination of all of the "extra bloat" which comes along with the support libs which 1.0 relies upon (that is no trivial amount, due to my packrat-like nature when it comes to utility code).

So, with our sights on portability, and also in the interest of a cleaner build process, the vast majority of the "support libs" have been factored either out or in. That is to say: some of the code (not much) got moved (back) in to s11n and the rest (the majority) was sent packing to CVS limbo. In any case, the s11n core tree is now 100% standalone, with some notes:

- The `zfstream` support lib is used by s11n if it is found, but it is not required. This is the *only one* of the 1.0 support libs which 1.1 now looks for - it no longer uses any of the others which 1.0 relies upon.
- All Serializers which ship with the library will be linked in with the main library, instead of as separate DLLs. This is primarily in the interest of easing portability to other platforms. Note that this does *not* change how the Serializers are *used* in client code, but affects how they are linked in with the main lib: they are still loaded via the dynamic-style interfaces (e.g. `s11nlite::create_serializer('MySerializer')`). If you are the only other person on the planet who actually does dynamically load Serializer DLLs this change will affect you, but if you're doing that then you know what needs to be done to fix it.

3.3 Factory code reimplemented

While the older factory/classloading code (named `cllite`) is functionally okay, and provides an adequate interface, its code base contains a lot of "evolution cruft". In December, 2004, i was offered a spot on the `pclasses.com` team, to assist them in their 2.x rewrite. The first assignment was to implement a new factory, which i did by taking the learnings from their 1.x factory, s11n's `cllite`, and some other experimental code. After it proved its worth in the `P::Classes` tree, i ported a copy into s11n. The newer factory is not markedly improved, functionally, but provides a more focused factory interface than `cllite` and has a couple new tricks to try out.

(It may be interesting to know that P 2.x has its own integrated copy of `libs11n`. *That's* why i want the s11n code to be easily forkable!)

3.4 `node_traits<>` changes, `s11n::data_node` replaced with `s11n::s11n_node`

To make a long story very short: the `data_node` type was "the original" abstract s11n container¹⁹, introduced in s11n 0.7.0. When the type traits system came along (version 0.9.3), i refactored `data_node` into a slightly more focused API, `s11n_node`. That class has been around since the summer of 2004, but hasn't been actively used within the s11n tree (only for testing the `node_traits`-related features). As of 1.1.0, `data_node` has been completely removed and replaced with `s11n_node`. Also, `s11n_node`'s API has changed slightly, to make it a bit leaner. Sorry for not having a deprecation period, but making the switch is actually much less painful than it sounds - even trivial (or a no-op) for most client-side code.

What this means for client code:

¹⁹ Not to misrepresent: i mean "the original" as in "the first one to exist in `libs11n`." The basic model for such containers had been demonstrated as early as summer 2000 in Rusty Ballinger's `libFunUtil`, if not also in other places, and was used, but in a much different way, in s11n 0.6.x and earlier.

- Users of the `s11n_lite::node_type` typedef normally simply need a recompile (which they would need anyway, because 1.1 is not binary-compatible with 1.0).
- Users of `node_traits<>` iterator-related typedefs and functions will need some slight modifications: don't use the (missing) typedefs, but go through the appropriate sub-typedef, so to say. For example: `node_traits<>::begin()` / `end()` and `node_traits<>::[const_]iterator` are now `node_traits<>::property_map_type::members` (they always were, but the "convenience" interface was removed because it was confusing to remember if it referred to the properties or the children).
- Clients who explicitly used `data_node` should globally replace that with `s11n_node`. This transition will normally be seamless *if you use* `node_traits<NodeType>` to manipulate your nodes (that is *The One and True Way*), otherwise other changes *might* be required to accommodate the API differences between the two node types. The APIs are functionally identical, but are *intentionally different* so as to trigger errors in the s11n core code if it does not hold to "the `node_traits<>` rule." (That is, the two node types have different APIs to *force me* to fix any s11n core code which isn't using `node_traits<>`!)
- Due to the above change, the `data_node.hpp` header of course no longer exists.
- These changes should not affect data files at all, because the two node types are fundamentally the same (only one string identifier in their output is different, but it's not significant for client purposes).

Users who follow the documentation and use `node_traits<NodeType>` to query and manipulate their data nodes, and clients who use template-defined Node Types rather than hard-coded ones, are mostly not affected by this change but may need to make some header-related fixes and a couple typename fixes. e.g. see the notes about some typedef-related changes and the removal of the `begin()` and `end()` members of `node_traits<>`. Their existence was logically ambiguous, with children and properties both competing for iterator types, and was confusing to remember which iterator `begin()` really returned. `node_traits<>` still contains all of the typedefs and accessors needed to get at that data, but the user will have to go one typedef or function call deeper to get it (but the client code's intention will also be clear to humans, which was not the case before without an additional lookup in the API docs).

3.5 New header conventions, faster compile times

Largely in the interest of bringing some sanity to the s11n build tree, and partly because i have an insatiable urge to hack build processes²⁰, we have undergone some significant build tree and header reorgs. Again. Yes, i know that's twice... er... three times in the past 12-month period. Learn to think of it "improvement via natural selection" and it doesn't hurt quite so badly. If it makes you feel any better (it does me), the very basic tests i have run show a cut in compile time by as much as 80%. That is, as much as 5 times faster compared to equivalent 1.0 code. Most client-side code will probably see compile times cut by 50%-70%, at least as far as the s11n-side of the compiles goes, and some code won't see much of a difference.

First off, the main Serializable registration header has been renamed: `reg_serializable_traits.hpp` is now called `reg_s11n_traits.hpp`, because that's what the file does - registers `s11n_traits<>`-related code.

Secondly, many headers have been renamed or consolidated into other headers (this mainly affects the i/o and proxy code, but also some of the core algorithms and functors).

The most notable reorg is how the serialization proxies for PODs and STL containers are registered. In 1.0 they were registered *en masse* via headers which included support for multiple containers. This is all fine and good, from an ease-of-use standpoint, but causes measurable (and human-noticable) increases in client-side compile times even for cases where most of the proxies aren't used. In an attempt to decrease client-side compile times, each proxy type now has its own header. All such headers follow common naming conventions and live in a new header subdirectory:

```
// register std::vector<T> proxy:
#include <s11n.net/s11n/proxy/std/vector.hpp>
// promote 'int' to a first-class Serializable:
#include <s11n.net/s11n/proxy/pod/int.hpp>
// algos and base proxies for list-like types, but no registrations:
#include <s11n.net/s11n/proxy/listish.hpp>
// algos and base proxies for map-like types, but no registrations:
#include <s11n.net/s11n/proxy/mapish.hpp>
... and so on...
```

The end effect is that clients must individually choose which proxies they will need. This is slightly unfortunate, but is a one-time cost of including the proper header(s). The main benefit is, for the vast majority of client-side

²⁰ Shameless plug: <http://toc.sourceforge.net>

cases, improved compile times. Even in the *worst* cases, compile times should be faster than with 1.0.x because 1.0 tries to install a lot of proxies which are almost never used. If this change *really* annoys users, they may make their own "mass-include" files and include all the proxies they want to. In fact, if compile times are not a concern to you, either because you are extremely patient or because you have access to the lab's Monster Computer, i recommend the mass-include approach, *but only for the sake of ease-of-use* when it comes to figuring out what proxies you need. For standard PC users, i don't recommend the mass-include approach at all, at least not unless you are unusually patient while waiting for your code to compile.

i have *attempted* to structure the proxy headers in a maintainable and extendable manner, such that it shouldn't take too much effort to locate the proper proxy header one needs, nor to add new proxies by following the current conventions. If *you* have suggestions for a better layout, please feel free to get in touch! (But be aware that your suggestion might be used, which might of course mean more code reorgs. ;)

3.6 Fetching class names of Serializables

In one of those, "You utter moron! You should have done this nine months ago!" moments, the `s11n_traits<SerializableType>` interface has been extended to include one static function:

```
static std::string class_name( const serializable_type * HINT );
```

See the API docs, in `traits.hpp`, for full details, but briefly: this replaces all of the older `class_name<>` and `classname<>()` kludgery which has been around since s11n's earliest days (0.2.x or 0.3.x, i think). The end effect is the same, functionally, but this approach fits in cleanly with the rest of the API, whereas the older approach did not (i never did like the old way, but it was necessary for a long time). This approach also allows users of 3rd-party libraries like Qt to use polymorphism-friendly `BaseObjectType::className()` [or similar] member functions, whereas the older approach did not directly support that at this level of the s11n architecture.

Design note: i am not at all happy about *not* providing a default of 0 for the `HINT` argument. However, given the usage of `s11n_traits<>`, which is only "extended" via template specializations, i also do not like the idea of relying on all specializations to provide that 0 in their interfaces. Also, in the case that it ever becomes useful to make `s11n_traits<>` a virtual base class, `class_name()` might become a virtual function (i repeat: *that is theoretically possible*, not a concrete plan), and default parameter values in virtual functions make me queasy, techno-philosophically speaking.

3.7 Client-extendible s11nlite

One of the more interesting additions to 1.1 is a polymorphic class which provides the same API as s11nlite: `client_api<NodeType>`. This effectively allows users to have an s11nlite interface for custom Node Types or to add custom stream handlers to the s11nlite API. s11nlite has been refactored to be based off of this new class, such that clients are able to subclass it and provide their own class instance to s11nlite via a back-door-shared-instance-injection technique. This can be used, e.g. to provide network support on top of s11nlite using tools like the experimental code at <http://s11n.net/ps11n/> (that code was the primary inspiration for the new class). For example, network-aware extensions to s11nlite can be plugged in to arbitrary s11nlite clients without their code, or s11nlite, even requiring a recompile. If some other desperate coder out there adds, say, Oracle support, your s11nlite client code will be able to use it without explicitly having to know about it. Consider, too, that we can actually use factories to dynamically load arbitrary instances of the `client_api<>`. Weird, eh?

3.8 ~/.s11nlite config file removed

In 1.0, s11nlite saves its configuration when the library shuts down. While this is all fine and good for a system where only one app uses s11nlite, it causes interference when multiple apps share s11n. For example, when App A sets `s11nlite::serializer_class('MySerializer')`, App B is going to get that default the next time it starts unless it sets its own (which might then affect App C... *ad nauseum*). Thus we take the simple route and remove it. The only affect this has on clients is that they might want or need to set a default Serializers when their app starts up, using `s11nlite::serializer_class()`.

While the majority of s11n users use the library in only one source tree, i currently use it in no less than six projects, and have often experienced problems with each app imposing its own idea of a default file format on the other apps. So, like so many other dead-ends of evolution, `~/.s11nlite` is gone.

Since the s11nlite config object was never really advertised as a feature, it is thought (hoped) that this change does not affect any clients.

Note that the serialization of an application-wide config file is trivial, but that techniques like finding a user's home directory are platform-specific (even under Unix, `$HOME` is not *always* the user's home directory).

See section 18.4 for info about a new class which provides behaviour similar to the older `s11n_lite` config object.

3.9 Exceptions conventions

As of version 1.1, i've finally started seriously working on defining exception conventions for the framework. Newer code fixes all known potential leaks which could have happened in the face of exceptions in 1.0.x. Also, many algorithms can finally make some guarantees which weren't possible in 1.0. If you are a 1.0 user with no compelling reason to upgrade, *this is the compelling reason*. These fixes theoretically can't be backported into 1.0 without either a really significant effort or significant incompatibilities with other 1.0 releases, neither of which i'm up for.

See section 16 for details, and please feel free to make suggestions.

4 Core concepts

Users who want to *fully* understand `s11n` should read this section carefully - here we detail the major components of, and terms used within the context of, the `s11n` architecture. Understanding these is critical if one wants to *truly* understand how the library works. That said, a *lot* can be done client-side without understanding *anywhere near* all of the gory details: one can get quite far by simply copying example code!

4.1 Terms and Definitions

Below is a list of core terms used in this library. The bolded words within the definitions highlight other terms defined in this list, or denote particularly significant data types. This bolding is intended to help reinforce understanding of the relationships between the various elements of the `s11n` library.

Note that some terms here may have other meanings outside the context of this software, and those meanings are omitted for clarity and brevity - here we only concern ourselves with the definitions as they pertain to us as users of `s11n`.

- **s11n** - several meanings:
 - A short-hand form of the word "serialization", used in many contexts.
 - The literal name of this software.
 - Serialization as a computing domain.
 - Other, more context-specific, meanings.
- **Data Node** or **S11n Node (S-Node)** - a generic term for map-like types which store arbitrary key/value properties and child nodes, plus some meta-data (like type information for the stored data). They are structured in a tree-like fashion, DOM-style. In `s11n_lite` this role is played by the `s11n::s11n_node` type, and core `s11n` supports any node type which conforms to the `node_traits<NT>` conventions (see below).

Note that using a Data Node's API directly from client code is discouraged. Please prefer the API provided by `s11n::node_traits<DataNodeType>` instead, as described in section 6.1.

As of version 1.1.3, the term Data Node is being slowly phased out in favor of S11n Node (or S-Node), as that term fits in better with this library.
- **Node Traits** (`s11n::node_traits<NodeType>`)- an interface for interacting with **S11n Nodes**. Conceptually similar to the standard library's `char_traits<char_t>`. See section 6.1.
- **serializable** (with a *small* "s")- the property of being able to be saved and to restore state. For example, to allow persistent object states across application sessions, network connections, etc.
- **Serializable Type** or **Serializable** (with a *big* "S") - any type for which `s11n` recognizes a **Serializable Interface**, either implemented directly by the **Serializable** type or via a **Serialization Proxy**. **Serializables** save their state in **S-Nodes** during **serialization** and restore their state from **S-Nodes** during **deserialization**.
- **Serializable Traits** (`s11n::s11n_traits<SerializableType>`)- a type for encapsulating `s11n`-related information about a **Serializable Type**. See section 6.2.
- **Serializable Proxy** or **Serialization Proxy** - a functor (optionally two) which registers with `s11n` as being the handler for de/serialization of a given type. By extension, the *proxied* type is considered to be a full-fledged **Serializable**. All **de/serialize operations** `s11n` performs on behalf of the proxied type are delegated to the proxy type. This allows, amongst other things, transparent serialization of 3rd-party classes and drastically simplifies the serialization of containers.

Proxies are *not* **Serializables** - they are, more properly, the implementation for a **Serializable's serialization operators**. (Got that?)

- **serialization, to serialize** - several meanings:
 - To save the state of a **Serializable**. In this library that is accomplished by storing the state in an **SNode**, which is conceptually identical to storing a *copy of it* in an STL container.
 - To save an **SNode** to a data stream via a **Serializer**. Stream-based serialization is normally called "saving".
 - Several other subtle, context-specific meanings.
- **deserialization, to deserialize** - the converse of **serialize**:
 - To restore the state of a **Serializable**, presumably using data from an **S-Node**.
 - To load an **S-Node** from an input stream. Stream-related deserialization is normally called "loading".
- **de/serialization** or **de/serialize** - shorthand forms of "deserialization and/or serialization" and "deserialize and/or serialize."
- **Load/Save vs De/Serialize** - By s11n convention, the words "save" and "load" are used when dealing with streams or files, and "serialize" and "deserialize" are used when dealing with saving or restoring the state of a **Serializable** to or from an **S-Node**. Sometimes the words are used interchangeably and, while it is technically correct in many cases, such usage is considered "marginally ambiguous" in s11n.
- **Serializer** - a type responsible for converting **S-Nodes** to and from a specific grammar (i.e., a data format). For example, some Serializers use an XML dialect while others use custom formats. Theoretically, any data which can be structured in a DOM-like fashion (even if only via logical transformation) can be handled by Serializers. In s11n *Serializers* are also always *Deserializers* (at least logically, in terms of the API interface).
- **serialization operators, de/serialize()**, or **Serializable Interface** - generic names for a pair of de/serialize functions which **Serializables** and **Serializable Proxies** have, regardless of the actual names or argument types of the functions. Sometimes also used to refer to the de/serialize functions within other interfaces, such as core library's `de/serialize()` functions.
- **de/serialization operations** - generic terms encompassing any functions which trigger a chain of events which lead through the s11n de/serialization core (and presumably back). In plain English: `s11n::de/serialize<>()`, and related functions, fall into this category. If we need a really technical definition, this would be pretty close to correct: any operations which end up forwarding through the `s11n_api_marshaler<>` (SAM) internal interfaces (section 17).
- **Default Serializable Interface** - Serializables which implement both of their serialization operators as `operator()`, and which follow the conventions laid out in section 5, are said to implement the *Default Serializable Interface*. Types which do this do not need to tell s11n what their serialization interface looks like - we will be able to pick them up automatically.
- **Classloader** or **Factory** - an interface used to load object instances based on a lookup key, potentially including dynamic searches for new types (e.g., via DLLs). In s11n this lookup key is conventionally the string form of a class' name. Classloaders are used during deserialization to load the proper type for a given node (this is necessary in order to support polymorphic deserialization). The s11n classloader has support for loading classes from DLLs, but that feature is not covered much in these docs because its operation is transparent to the API. Classloaders work primarily not off of specific "concrete" types, but off of **Interface Types**, as described briefly below. For more detail than you probably want to know about these, see the summary paper at: http://s11n.net/papers/#classloading_cpp
- **T's classloader, or the T classloader** - Refers to the classloader (factory) which uses type T as its point of reference for registering and loading classes. More specifically, it (currently) means `s11n::fac::factory_mgr<T>`, though the exact factory implementation which s11n uses is not a defined part of the public interface. For proper polymorphic deserialization, subtypes of T should be registered with T's classloader, regardless of whether or not they also register with their own classloader (e.g. `factory_mgr<SomeTSubType>`).
- **Interface Type** [note: in s11n 1.0 these are referred to as Base Types, which is marginally incorrect and definitely more ambiguous than Interface Type.] - in s11n, especially in the context of a **classloader**, this is used to mean the base-most type which a given **classloader** "knows about." This type is used for registering subtypes of Interface Types with the Interface Type's **Serializable Interface**, and is critical for classloading purposes. In a broader sense, Interface Types are used as contexts for marshaling the s11n and client-side **Serializable Interfaces** into internally-compatible forms. The abstract topic of Interface Types is covered in more detail in a paper written as part of this project: http://s11n.net/papers/#classloading_cpp
- **Streamable [Types]** - In the context of s11n this means any type for which `ostream<<` and `istream>>` operators can be applied to successfully save and restore the state of an object of that type. This inherently includes all PODs, `std::string` (though with some caveats involving whitespace handling), and any client-supplied types which meet these conditions. This also implicitly *excludes* all pointer-qualified types (but note that s11n often handles objects of types (`SerializableT`) and (`SerializableT *`) equally). **Serializables** are *not* implicitly Streamable, as s11n does not deal with streams at its core, and thus the **Serializable** interface is stream-ignorant.

- **SAM, the s11n API Marshaler** - SAM is the layer of s11n responsible for acting as a communication channel between s11n's internal API and any client-side APIs, including, but not technically limited to, forwarding requests to **Serializable Proxies**. SAM allows clients to transparently proxy the s11n interfaces, as covered in section 17. Clients will almost never have to know about SAM, but it does play a significant role internally.
- **core s11n or the s11n core/kernel** - These are generic terms referring to the core-most functions in s11n. Specifically, this is limited to the classloader-related functions in the `s11n::cl` namespace, the `s11n::de/serialize()` variants, and `s11n::s11n_api_marshall<>`. Everything else, from the **Serializers** to the s11n-lite interface, is built around this *tiny* core.
- **POD** - Plain Old Data. In s11n this term does not have *quite* the same meaning as the C++ standard applies to it: we use it only to mean basic, built-in data types (`int`, `char`, `double`), plus `std::string`. In the C++ standard the term does not include `std::string` but includes `structs` which contain only PODs [CTM2005], and thus our usage includes a subset of the standard's definition. Common usage of the term does not include `structs`, so i don't feel bad about this slight mis-use of the term.

Did you get all that? Don't worry - you don't need to memorize this list, but if you find yourself confused by a term in this documentation, try looking it up in the list above.

Using the library is not as complex as the above list may imply, as the rest of this documentation will attempt to convince you. Yes, the details of serialization and classloading, especially in a lower-level language like C++, are *downright scary*. s11n tries to move the client as far away as possible from those scary details, and it goes to great pains to do so. However, some understanding of the above terms, and their inter-relationships, is critical *fully* understanding the library.

Some non-s11n-related terms show up often enough in this documentation that readers not familiar with them will be at a disadvantage in understanding the documentation. Briefly, they are:

- **i.e.** - "in other words" or "in effect" (from the Latin *id est*²¹).
- **e.g.** - "for example" or "example given" (from the Latin *exempli gratia*).
- **Algorithm** - we use the same general meaning as in common STL usage: a computation, normally one which is genericized in form such that it can be applied to a wide range of types which meet a published set of conventions for that algorithm. Like **functors**, understanding algorithms is essential to effectively using the STL, and the two often go hand-in-hand. For *numerous* well-published examples of algorithms see those in the STL itself, defined in the ISO-standard `<algorithm>` header file. s11n includes many serialization-related algorithms and functors.
- **Functor** - a function or a struct/class type implementing function-call semantics. i.e., a type implementing one or more `operator()` member functions. Functors are a cornerstone of all STL-style development, and must be well understood before one can make full use of s11n, or the STL for that matter.
- **ODR, the One Definition Rule** - C/C++'s rule which, put simply, basically states that no type may be *defined* (i.e., implemented) more than one time in any given binary or library. This is not an arbitrary rule, but a technological limitation, akin to `std::map` being able to contain no more than one object with a given lookup key. In any case, it's rather a *sane* behaviour, if you ask me. In s11n ODR is an oft-heard term because its template-based nature, in particular its use of macros and header files to generate "behind-the-scenes" utility and marshaler class template specializations at compile-time, makes it quite susceptible to ODR violations if some simple, non-obstructive rules are not followed (as described elsewhere in this manual). (Trust me, once you realize how it works this is never a practical hindrance, and it's trivial to avoid once you seen it happen it a few times and understand its nature.) With the release of version 0.8.0, all commonly-occurring ODR-related problems are believed to be solved. (i haven't personally seen an s11n-caused ODR violation since the 0.8.x series, except when i have incorrectly double-registered a proxy in the same source file.)
- **Style Points(SP)** - an abstract, often poorly-understood and underestimated, unit of measurement of "how much *Style*" a particular piece of code exhibits. Poorly-designed code gets minus points, whereas especially clever code may get plus points (or may, as is occasionally the case, actually be too clever for its own good, and get no points at all). The measurement system for Style Points is not standardized. One common way for one developer to communicating that s/he wishes to assign SP to, or subtract SP from, another developer is to say something like, "+1", or "-1". A phrase like, "cool code!" implicitly carries at least one SP, whereas the phrase, "great hack!" or "you rock!" is generally worth several SP (at least from the receiver's perspective). It is significant to keep in mind that SP declared by non-developers simply go to `/dev/null` - they neither count nor discount the recipient, except possibly in his or her own ego²². Additionally, the amount of SP a given reward or penalty gives or takes may be adjusted by the relative experience levels or reputations of the giver and receiver. e.g. a 6-month C++ newbie giving +1 SP to a 10-year

²¹ <http://www.wsu.edu:8080/~brians/errors/e.g.html>

²² And we programmers, by and large, have a reputation for living the majority of our lives in exactly that space. ;)

veteran is not worth *nearly* as much the other way around.

The giving of Style Points is sometimes referred to as "schenking" (past tense: *schanked* or *schenkt*), derived from the German verb *schchenken*, meaning "to give [free of cost/as a gift]."

As software developers mature²³ they invariably begin, at some indefinite point, to concentrate on Style as much as they do on the nature of the algorithms they develop. This is a natural part of a developer's growth as a professional, just as it is in any field, and thus experienced coders can generally "pick up SP" much more readily than greenhorns can.

4.2 The Official *Grossly Oversimplified Overview* of the s11n architecture

"Like your scrotum, here it is in a nutshell..."

Bloodhound Gang (the band, not the TV show or children's books)

s11n is built out of several quasi-independent sub-modules. "Quasi-independent" meaning that they mostly rely on *conventions* developed within other modules, but not necessarily on the exact *types* used by those modules. Such design techniques are a cornerstone of templates-based development, and will be a well-understood principal to STL coders, thus we won't even begin to touch on its benefits, uses, and multitudinous implications here.

*Shameless Plug*²⁴:

This particular aspect of s11n's design is critical to s11n's flexibility, and is one of the implementation details which catapults it *far* ahead of traditional serialization libraries. It is this aspect which allows, for example, client libraries to transparently adapt this framework's interfaces to the client's interface(s), and to transparently adapt *other clients'* Serializable interfaces (and, additionally, transparently adapt to *them*). In most other libraries this model is the other way around: the client has to do all adapting himself. Consider, e.g. that *any* type can be converted to a Serializable without, e.g. subclassing anything at all. That is, a client can have 1047 different classes - each with their own serialization interfaces - and they can all transparently de/serialize each other *as if they all had the same function-level interface*²⁵.

Enough plugging. Let's briefly go over s11n's major components, in no particular order:

- **ClassLoader** - a factory for creating classes based on lookup keys (e.g. class name). This is a critical element for proper polymorphic deserialization, particularly when loading classes on-the-fly from external sources (e.g. a DLL).
- **s11n::s11n_node** - this is the reference implementation for the **S11n Node** (a.k.a. Data Node) concept. This is supported by all of supplied node-related algorithms and functors, though they actually have no direct dependencies on it. It is considered poor style use call the Data Node API directly from client code - using the **s11n::node_traits<NodeType>** interface is *highly* preferred, for compatibility with 3rd-party node types and for future compatibility with new node types. For example, s11n 1.0.x uses a different node type (**s11n::data_node**), and using the **node_traits<>** to access nodes makes this transition transparent.
- **Core de/serialize() functions** - a set of functions which hide the API marshaling that goes on for translating arbitrary Serializable interfaces into something each other can use. At the application level, these functions typically make up the heart of the client-side s11n interface, whereas at the library- and class- levels the available functors and algorithms are much more likely to play a heavy role. It may be interesting to note that the core API is made up of less than 50 lines of code.
- **Serialization API marshaler (SAM)** - the core de/serialize functions pass all of their request through this internal layer. These types can be swapped out transparently, customizing the serialization interface on a per-base-type case. This feature is used, for example, to direct serialization through Serializable Proxies, or to implement pointer-to-reference type translation as needed. These marshalers filter every single de/serialize call made via the core, and thus the ability to replace them on the client side gives client code 100% plug-in access to the framework's de/serialization core, without having to know the details of how everything is marshaled. SAMs can then do almost whatever they like with the API, except change parameter constness for nodes and serializables - they may add arguments as they wish! This can be used, e.g. to implement framework-enforced data versioning. **SAM** is covered in section 17.

²³ As developers, of course, not necessarily as human beings.

²⁴ Such a plug is typically worth approximately *-1 Style Point*, a cost from which this plug is not exempt. In fact, these docs have so many shameless plugs and outbursts of jubileum that I'll go ahead and dock the document as a whole *-10 SP*. ;) (I wouldn't be preaching if I didn't believe honestly it, though, so the devotion's gotta be worth a couple of SP!) *What a Style Point?* See section 4.1.

²⁵ Whereas they do all implicitly share a common *logical* interface - that of a Serializable, as defined by s11n's conventions.

- **Type Traits** (section 6) - as of version 0.9.3 these types are used to encapsulate interface information for Serializables and Data Nodes. Users of the STL may be familiar with standard traits types such as `iterator_traits<>` or `char_traits<>`. The s11n traits types, `s11n_traits<>` and `node_traits<>`, play similar roles as those types do in the STL. Note that `s11n_traits<>` and SAM overlap in some ways, as described in section 17.
- **Serializers** - these objects are responsible for marshaling S-Nodes to and from specific file formats (also known as *grammars*). The library currently s11n ships with several Serializers, supporting a variety of data formats. All Serializers shipped with s11n are available to s11n-lite, but s11n-lite restricts itself, for purposes of *saving* data, to one of them (*which one* it uses is not strictly defined by the interface, and may easily be defined by the user). s11n-lite does not need to be told what format to use for loading, as that is determined dynamically (see sections 14.1.1 and 14.1.4).
- **s11n-lite** - a tidy little interface providing a wrapper around the above layers, providing for most common client object serialization needs. Intended also as a sample client-side interface implementation. That is, by implementing something like s11n-lite a project can completely hide its objects from any direct knowledge of libs11n, helping to support the "non-intrusion principal" which s11n works hard to uphold. For an example of this, see the P::SIO module in the P::Classes 2.x source tree (via <http://pclasses.com>), where we have implemented a custom s11n-lite-like interface to suit the needs of that project better.
- Generic helper functors and algorithms to support internal and client-side manipulation of Data Nodes and Serializers, also helpful for s11n-lite.

There are also a number of less-visible support layers/classes/functions. See the README file for an overview of where each part of the library lives in the source tree. The API docs reveal the whole spectrum of available objects (many of which are internal or special-case, and can be ignored by clients).

Some of the sub-sub layers exist purely as code generated by macros (such as the classloader registration macros), e.g. to install client-specific preferences into the library at compile-time.

4.3 Process Overview

4.3.1 Serialization

In the abstract, this is normally what happens for a serialization operation:

1. Client requests the serialization of a Serializable. This is initialized by passing the Serializable into a data container (e.g. an S-Node) via the s11n serialization interface (e.g. `s11n-lite::serialize()`).
2. s11n proxies the request to the registered Serializable Interface and passes the target S-Node and source Serializable to the registered interface.
3. The serialize operator's implementation should save the Serializable's state into the data node. It returns true on success and on error returns false or throws an exception.
4. s11n returns a data node to the client, presumably populated with the data from the Serializable.
5. Client selects a Serializer type and sends the Node to it, along with a destination stream/file.
6. Serializer formats the Node into the Serializer's grammar.
7. The client gets notification of success or failure (true or false, respectively, or potentially an exception).

Recursive serialization can be triggered, e.g. in a serialization operator's implementation where a child Serializable is serialized.

Note that in s11n-lite the Serializer selection steps are abstracted away to simplify the interface.

4.3.2 Deserialization

A client-initiated deserialization request in s11n normally looks more or less like this:

1. Client requests the deserialization of a Serializable object from a data stream/file.
2. s11n analyses the stream to find a matching Serializer class, then passes the stream off to the that class.
3. The Serializer parses the stream into a tree of S-Nodes and returns the root node to s11n. Obviously, if there is no Node then processing stops here with an error (typically, false or 0 is returned, though an exception may also be thrown).
4. s11n looks at the root Node to determine which Serializable Type to instantiate. If it fails to find the class, or cannot instantiate the requested type, processing stops with an error (typically false or 0 is returned, though an exception may be thrown).
5. s11n marshals the data-to-be-deserialized to the registered (De)serialization Interface for Serializable's type.
6. Deserialize operator's implementation should restore the Serializable's state from the source Data

Node. If it returns false or throws then processing stops. In the case of an error it may do post-error cleanup on the object to prevent leaks of resources allocated during deserialization.

7. `s11n` destroys the now-unnecessary S-Node tree.
8. `s11n` returns a (Serializable *) to the client, which the client now owns.

The interface also supports deserializing nodes directly into arbitrary Serializables, effectively bypassing the first four of the above steps and not returning a pointer to a new object (it uses the target object the user gives it). Also, clients may stop at point 7 if they are only interested in the raw data, as opposed to wanting the objects the data represent. For example, the `s11nconvert` and `s11nbrowser` applications (sections 21.1 and 21.2) never rely on a specific Serializable Types, and only work with S-Node trees.

4.4 Node Names and Property Key naming conventions (IMPORTANT!)

When saving data each node is given a name, fetchable via `node_traits<NodeType>::name()`. Node names can be thought of as property keys, with the node's content representing the value of that key. Unlike property keys, node names need not be unique within any given data tree. All nodes have a default name, but the default name is not defined (i.e., clients can safely rely on new nodes having *some* Serializer-parseable name).

In terms of the core `s11n` framework, the key/node names client code uses are irrelevant, but most data formats will require that they follow the syntax conventionally used by XML nodes and in most programming languages:

Alphanumeric and underscores only, starting with a letter or underscore.

Any other keys or node names will almost certainly *not be loadable* (they will probably be saveable, but the data will be effectively corrupted). More precisely, this depends on the data format you've chosen (some don't care so much about this detail).

Numeric *property keys* are another topic altogether. Strictly speaking, they are not portable to all parsers. More specifically, numeric keys (even floating-point) are handled by most of the parsers supplied with this library (even `funxml` and `simplexml`, but not `expatxml`), but the data won't be portable to more standards-compliant parsers. Thus, if data portability is a concern, avoid numeric property keys and node names altogether.

Serializable classes normally do not need to deal with a node's `name()` except to de/serialize child Serializables. There are many cases where client code needs to set a node name manually, but these should become clear to the coder as they arise.

4.5 Overview of things to understand about s11n

After reading over the basic library conventions, users should read through the following to get an overview of what topics which should be understood by clients in order to effectively use the `s11n` framework. Much of it is over-simplified here - this is an overview, after all. Additionally, some of it is true for `s11n-lite`, but only partially true for core `s11n`.

- S-Nodes are the basic types used to store arbitrary key/value pairs and child objects. They follow a DOM-style interface, so their usage is fairly straightforward. The core library and generic algorithms support any S-Node type which can be proxied via `s11n::node_traits<>` (section 6.1).
- The entire client-side interface for loading and saving all objects is declared in `<s11n.net/s11n/s11n-lite.hpp>`, in the `s11n-lite` namespace. The core code, and many node-related functors and algorithms are available in these namespaces: `s11n`, `s11n::list`, `s11n::map`, `s11n::va`. That said, clients may directly use the core `s11n`, bypassing `s11n-lite` completely, but using `s11n-lite` is *highly* recommended.
- `s11n` is very container/functor/algorithm based, so its usage should be familiar to experienced C++ users (especially users of the STL).
- `s11n` does not enforce a specific Serializable interface, but inherently supports the so-called *Default Serializable Interface*. Client-side classes which implement the default Serializable interface (described later) need no special registration as being Serializable types. Custom interfaces and proxies are *easy* to install, as described later.
- `s11n`'s core is not stream-oriented, but container-oriented. That is, we serialize data to and from containers, and those containers get formatted to (or from) streams by Serializers. Thus `s11n` doesn't really care about file formats - its core interface is 100% data format agnostic. For saving, clients must declare a format, but loading is dispatched to the appropriate parser depending on the content of the stream. That said, `s11n-lite` uses a default Serializer, so clients who don't care about the underlying data format need never worry about this highly overrated detail.
- Classloaders and their "InterfaceType" types are important concepts to understand in `s11n`, mainly for template-types reasons. They are covered *in detail* in the classloader documentation, and will be

explained a bit later on. *All* types which are to be *deserializable* must be registered with an "appropriate classloader." What that *really* means, in all its technical glory, could easily turn into whole document! Be assured that this doc will try to tell you what you need to know in order to register your classes (it is 100% non-intrusive on classes). The hope is that most s11n use cases won't require much understanding of the subtleties of the classloader framework.

4.6 Notes on error/success values (i.e., justifying the bool)

s11n uses, almost exclusively, bool values to report success or failure for de/serialize operations. The reasons that bool was chosen are detailed, but here's a summary:

- SOME error value is needed. Integer values must either be mapped to a known set of error codes or be interpreted client-dependently. Neither of those approaches are terribly suitable for s11n, largely due to its inherently abstract and generic nature.
- Based on usage history, i felt it was unnecessary to employ exceptions as the standard means of error reporting. (i partially regret this, but still generally feel that imposing exception conventions on the clients would not be a good idea.)
- If we consider the standard ostream<< and >>istream operators for a moment: yes, it is technically possible to check for an error after an extraction/insertion by checking the stream's state, but in practice this is *almost never* done, at least for ostream. Thus, i/ostream error checking conventions are oddly similar to s11n's, probably due to their logically similar roles as i/o marshalers.
- Related to the previous point: s11n's core is container-based, and how many coders check for proper insertion after a `push_back()` or `insert()`? None, because those operations (perhaps only by convention?) simply do not fail.
- i actually knew a coder once who (in Java) chose to return the String "success" to indicate success and non-"success" to indicate failure. i figure that's also not appropriate for s11n. ;)

s11n's conceptual ancestor, Rusty Ballinger's libFunUtil, uses `void` returns for its de/serialize operations, which means that clients essentially can't know if a de/serialize fails. When designing s11n i strongly felt that clients need at least add *some* basic level of error detection, and finally settled on plain old booleans. There is in fact a comic irony in that decision: it is so rare that a de/serialization fails, that a void return type would do just as well for 99% of cases!

The seeming shortage of de/serialization failures can primarily be attributed to the following:

- The vast majority of the client-side part of s11n doesn't work with i/o streams (in particular, with files).
- The points at which Serializables are given data nodes are far away (in interface terms) from the stream operations. Stream operations are, *by far*, the most likely point of failure in a serialization library (bad input format, file does not exist, out of disk space, write access fails, NFS connection cut, blah blah yada yada yada).
- The s11n core is container-based, and container insertions and extractions, as a general rule, do not fail. Also, container searches only fail in the sense that the sought-after data simply isn't there.
- In the event of a stream- or grammar-level input failure the process will fail early enough that no deserialize operators are be called, so they can't very well fail, can they?

[... much later ...]

While returning a bool for a single de/serialization operation still seems reasonable, the logic behind it rather breaks down when a *tree* of objects is serialized. If any given object returns false the the serialization as a *whole* will fail. This implies that whole trees can be spoiled by one bad apple (no pun intended). In a best-case scenario only one branch of the tree would be invalidated, but... *is that a good thing*, to have partial data saved/loaded and have it flagged as a success? Of course not, thus s11n must generally consider one serialization failure in a chain of calls to be a *total* failure. This is its general policy, though client/helper code is not required by s11n to enforce such a convention²⁶.

Furthermore, some specific operations, such as using `std::for_each()` to serialize a list of Serializables, may [will] have unpredictable results in the face of a serialization failure. Consider: in that case there is no reasonable way to know which child failed serialization, as `for_each()` will return the overall result of the operation. If the functor performing the serialization continues after the first error it will produce much different (but not necessarily more valid) results than if it rejects all requests after a serialization failure. The `subnode_serialize_f<>` class, for example, refuses to serialize further children after the first failure, but this is purely that class' convention, not a rule.

Ah... there is no 100% satisfying solution, and bools seem to meet the middle ground fairly well.

[... much later ...]

²⁶ Especially when s11n's author cannot even decide if s11n currently does The Right[est] Thing ;). It's mainly a philosophical question at this point, and those are often the most difficult ones in software design. :/

As of version 1.1 we've introduced proper exception handling: more info about this is in section 16.

4.7 s11n and Patterns

"Patterns" is a term we've all come to know and love over the past decade. While i am no Pattern Guru, and cannot name more than a couple off the top of my head, i thought it might be interesting to list the major components of the library and the Patterns they [would seem to] follow. This might help some users understand the library somewhat better...

4.7.1 The core

The core of the library is essentially a *Proxy*. All that it does is use templates to select types, and then call a known interface in that type, passing on the caller's arguments and returning the same value as the proxy.

4.7.2 Classloader

The classloading layer is, quite naturally, a *Factory*: it maintains a mapping of keys to functions which return new objects.

4.7.3 Proxies

Proxies are, quite non-intuitively, normally more like *Visitors* than *Proxies*. This really depends on the implementation, but in practice most are *Visitors*. The original design goal of the s11n proxies was to do only API marshaling (proxying), but it quickly became clear that they could do much more than that. By that time, though, the term Proxy was already in use and there was no reason (at the time) to think it wasn't appropriate.

Proxies normally implement one of three approaches:

1. They simply pass on their arguments to a *known* Serialization Operator in the Serializable type they proxy. In this sense they are naturally *Proxies*.
2. They implement the de/serialization logic for a Serializable type. In this sense they could be considered *Visitors*.
3. They pass on all arguments/return values to/from *algorithms* which perform #2. Again, in this sense they are *Proxies*.

For you Pattern Gurus out there: is there a separate Pattern for *API Marshaler*, or is that just a fancy word for Proxy?

4.7.4 i/o

The i/o layer is conceptually very similar to the proxying layer, though with much less indirection going on. This layer would appear to be mainly a *Visitor*, at least for output purposes, but there might be closer Pattern matches, so to say. In some sense it is also a Factory of S11n Nodes.

4.7.5 s11nlite

s11nlite is a classic *Wrapper*, which probably also falls into the category of *Proxy* or *Marshaler*.

5 Serializable Interfaces: overview and conventions

Rather than overload you with the details of this right up front, we're going to *grossly oversimplify* here - to the point where we're almost lying - and tell you that the following is *the* interface which s11n expects from your Serializable types.

Each Serializable type must implement the following two methods:

A **serialize operator**:

```
[virtual] bool operator() ( NodeType & dest ) const;
```

A **deserialize operator**:

```
[virtual] bool operator() ( const NodeType & src );
```

It is important to remember that *NodeType* is actually an abstract description: any type meeting s11n's S-Node conventions will do. s11nlite uses, unsurprisingly, `s11n::s11n_node` as the reference implementation for the

NodeType concept.

The astute reader may have noticed that the above two functions have the same signature... *almost*. Their constness is different, and C++ is smart enough to differentiate. The s11n interface is designed such that it is very difficult for clients to have an environment where ambiguity is possible. The justifications for the constness are essentially this: when serializing, clients need to know that the objects they are saving are not changed during the process. Likewise, when loading, an object needs to be able to change its state to match that of the loaded data.

These operators need not be virtual, but they may be so. Serializable proxy functors, in particular, are known for having non-virtual serialization operators, as are, of course, monomorphic Serializable types.

The truth is that s11n only requires that the argument be a compatible data node type and that the constness matches. s11n's core doesn't care what function it calls, as long as you tell it which one to use - how to tell s11n that is explained in section 12.

Trivia: When the de/serialize operators are implemented in terms of `operator()`, with the above-shown signatures, a type is said to conform to the Default Serializable Interface.

5.1 Serialize Operator conventions

- If the type is polymorphic, it **must** set its class name in the node, e.g. using `node_traits<NodeType>::class_name()`. This is currently the only 100% reliable way to get the proper class names of your Serializable subtypes for use during during deserialization. (This is made clearer later via examples.) Monomorphic types can be reliably given a name by the framework, and normally no class name needs to be called for them (SAM does this - section 17, and proxies sometimes re-set it). If this operator calls a parent type's serialization operator, the class name should be set *after* calling the inherited operator, such that the *subtype's* class name is stored.
- Should save the object's state to the destination node, presumably using the destination's public API and the s11n functors/algorithms designed for such operations. State-saving may continue recursively for Serializable child objects.
- Returns true on success, false on error. May throw or propagate arbitrary exceptions.

5.2 Deserialize Operator conventions

- Should restore the state of an object via the node it is given, plus any sub-nodes, if needed. State restoration may continue recursively for collecting Serializable child objects.
- The core library generally makes sure that nodes are passed to objects of the types which serialized the nodes, but users may "brute-force" any node into any Serializable if they wish to. It is not the job of the deserialize operator to check that it has received a node for the proper type. It may do so, if it wishes, but this is out of line with s11n conventions, and not recommended.
- The core library only calls the deserialize operator *one time per object*, but it is possible that client code will trigger it multiple times for a given object. Thus any lists, pointers and whatnot should be cleaned up before restoring an object's state, to avoid leaking resources or duplicating container entries. More information about this can be found in section 19.
- Returns true on success, false on error. May throw or propagate arbitrary exceptions. If it throws, it should ensure that it does not leak any resources allocated as a side-effect of deserialization, including resources allocated by recursive deserialization. (This is not as difficult as it sounds: see section 16.)

5.3 Data Node class names (IMPORTANT!)

Let us repeat this many times:

```
while( ! this->gets_the_point() )
    std::cout << "The importance of class_name() in the s11n framework cannot be understated.\n";
```

(Don't be ashamed if your loop runs a little longer than average. It's a learning process.)

`class_name()` is part of the `node_traits` interface, and is used for getting and setting the class name of the *type* of object a node's data represents. This class name is stored in the meta-data of a node and is used for classloading the proper implementation types during deserialization. By *convention* the `class_name()` is the string version of the C++ class name, including any namespace part but minus any qualifiers like pointeriness and template parameters, e.g. "foo::bar::MyClass". The library does not enforce this convention,

and there are indeed cases where using aliases can simplify things or make them more flexible. See the classloader documentation for hints on what aliasing can potentially do for you.

Client code *must*, unfortunately, call `class_name()`, but the rules are very simple:

- Serializables (or their proxies) must set the target node's `class_name()` in their *serialize operator* (not the deserialize operator), passing it the string name which the client code will later expect to be able to load the class with. When using the default Serializable registration techniques, you should pass the class name defined in the `S11N_TYPE_NAME` macro passed in to the registration supermacro (section 13.1.3).
- If a Serializable class inherits serializable behaviour from a parent type, the subclass must set `class_name()` *after* calling the parent's implementation of the Serialize Operator to ensure the proper subclass type gets into the node. Also, if the parent's operation fails, the child should normally immediately return false.

Some algorithms parse data directly from data nodes, irrespective of the node's `class_name()`, and this is perfectly kosher. One example is the `de/serialize_streamable_xxx()` family of functions: they use "raw" data nodes, to avoid a number of problems involved with registering proper class *names* for arbitrary containers' classloaders.

For more on class names, including how to set them in a uniform way for arbitrary types, see section



5.3.1 Example of setting a node's class name

Here's a sample which shows you all you need to know about the bastard child of the s11n framework, `class_name()`:

Assume class A is a Serializable Interface Type using the *Default Serializable Interface* and B is a subtype of A. In A's *serialize* (not *DEserialize*) operator we must write:

```
s11n::node_traits<SNodeType>::class_name( node, 'A' );
```

In B's we should do:

```
if( ! this->A::operator()( node ) )27 return false;  
s11n::node_traits<SNodeType>::class_name( node, 'B' );
```

It is not strictly necessary that a subtype return false if the parent type fails to serialize, but it is a good idea unless the subtype knows how to detect and recover from the problem.

Follow those simple rules and all will be well when it comes to loading the proper type at deserialization time²⁸. To extend the above example, after the node contains B's state, we can do this:

```
A * a = s11n::deserialize<A>( node );
```

(Note that we call `deserialize<A>()` with A because that's the Interface Type which registered with s11n.)

That creates a (B*) and deserializes it using B's interface. Why? Because node's `class_name()` is "B", and the A classloader will load a B object when asked to (assuming it can find B - if it cannot it will return zero/null, or possibly throw).

Let's quickly look at two similar variants on the above which are generally *not* correct:

```
B * a = s11n::deserialize<A>( node );
```

That won't work because there is no implicit conversion possible from A to B. It will fail at *compile time*. That one is straightforward, but the details for this one are fairly intricate:

```
B * a = s11n::deserialize<B>( node );
```

This will not fail to compile, but will *probably* not do what was expected. In this example B is now the Interface Type for classloading/deserialization purposes, which has subtle-yet-significant side-effects. For example, if B is never registered with *the B classloader* then the user will probably be surprised when the above returns 0 instead of a new, freshly-deserialized object. If B is indeed registered with B's classloader, and B (as a standalone type) is recognized as a Serializable, then that call would work as expected: it would return a

²⁷ See section 5.4 for why you should never directly call a Serializable's serialization API. This particular case is one of two which simply cannot be avoided.

²⁸ That is, assuming the subtypes are properly registered with the classloader.

deserialized (B*).

5.3.2 Using local library support for `className()`

Some heavily object-oriented libraries, like Qt (www.trolltech.com), support a polymorphic `className()` function, or similar, to fetch the proper, polymorphic class name of an object. If your trees support this, *take advantage of it*: set the node's class name one time in the base serialization algorithm (your proxy or the base-most implementation of your hierarchy) if you can get away with it! The sad news is, however, that the vast majority of us mortals must get by with doing this one part the hard way. :/ There are actually interesting macro/template-based ways to catch this for monomorphic types, but no 100% reliable way to catch them for polymorphs has yet been discovered. (*Hear my cries, oh mighty C++ Standardization Committee!*)

This approach is demonstrated in the s11n sample source code, in `src/client/sample/classname.cpp`.

5.4 Cooperating with other Serializable interfaces

Despite common coding practice, and perhaps even common sense, client Serializables *should not* - for reasons of form and code reusability - call their own interfaces' `de/serialize` functions directly! Instead they should use the various `de/serialize()` functions. This ensures that interface translation can be done by s11n, allowing Serializables of different ancestries and interfaces to transparently interoperate. It also helps keep your code more portable to being used in other projects which support s11n. There are *exactly three* known cases where a client Serializable must call its direct ancestor's `de/serialize` methods directly, as opposed to through a proxy. The first two are calling the parent implementation in their `serialize` and `deserialize` implementations. In those two cases it's perfectly acceptable to do so, and in fact could not be done any other way. The final case is when you want or need to bypass the internal API marshalling. Any other usage can be considered "poor form" and "unportable." If you find yourself directly calling a Serializable's `de/serialize` methods, see if you can do it via the core API instead (tip: *you probably can*²⁹).

For example, instead of using this:

```
myserializable->serialize( my_data_node );  
// NO! Poor form! Unportable!
```

use one of these:

```
s11nlite::serialize( my_data_node, myserializable );  
// YES! Friendly and portable!  
s11n::serialize( my_data_node, myserializable );  
// Fine!
```

Note that there are extremely subtle differences in the calling of the previous two functions: the exact template arguments they take are different. In the case of monomorphic types C++'s automatic argument-to-template type resolution suffices to select the proper types, so specifying them via `serialize<X>` syntax is unnecessary. When serializing monomorphs, being explicit should never be required. When using polymorphs, it may be necessary to explicitly give the base-most (interface) type, so that the subtype's type is not accidentally selected (which will lead to no good). It is always safe to do so, in any case, and s11n's author encourages always being explicit in this regard, to avoid potential confusion or subtle errors downstream.

In terms of Style Points (section 4.1), calling a Serializable's API directly, except where specifically necessary, is immediately worth a good -1 SP or more, and may forever blemish one's reputation as a generic coder. To be perfectly clear, though, calling the local APIs directly *does not* have any direct effect on s11n. This convention is primarily to help ensure portability of serialization functionality between disparate s11n-enabled types.

5.5 Member template functions as serialization operators

If a Serializable type implements template-based serialization operators, e.g.:

```
//serialize operator:  
template <typename NodeType>  
bool operator()( NodeType & dest ) const;  
  
// deserialize operator:  
template <typename NodeType>
```

²⁹ Alas, unless, you have some unusual needs, e.g. you need customized recursive `de/serialization` to go around the internal marshalling process.

```
bool operator()( const NodeType & src );
```

and they use the `s11n::node_traits<NodeType>` interface to query and manipulate the nodes, then their `Serialize` methods will support any `NodeType` supported by `s11n`. Note that `s11n-lite` hides the abstractness of the `NodeType`, so users wishing to do this will have to work more with the core functions (which essentially only means using `NodeType` a *lot* more, e.g. `functionname<NodeType...>(...)`).

Using member template functions has other implications, and should be well-thought-out before it is implemented:

- May require including (no pun intended) the implementation code in the header file.
- Compilers do not completely check template functions until they are called, so there might be a compile-error-in-waiting as coders tweak bits without testing them (what, me? ;).
- Member template functions cannot be virtual. (This is a C++ restriction, not `s11n`-imposed.)

Despite those seeming limitations, experience suggests more and more that templated de/serialize operators generally offers more flexibility than non-templated. In the case of monomorphic types and proxies, there is almost never a reason to *not* make these operators member templates, and there are several good reasons to do so:

- The class can work with any Data Node type, instead of just, e.g. `s11n-lite::node_type`.
- This is the only known effective way to proxy requests for class templates, e.g. STL containers, as it allows a single pair of functions to handle de/serialization for a whole family of types. e.g. two functions which can handle `list<int>`, `list<double>`, `list<char>` ...
- Common C++ literature suggests that smart compilers can eliminate at least some of the middle-man code in many common functor-related constructs.

6 Type Traits

In version 0.9.3 a Type Traits-based system was added to the framework to encapsulate information about the `S11nNode` and `Serializable` interfaces.

The traits types live in the namespace `s11n` and are declared in the file `traits.hpp`.

In short, the traits types encapsulate information about `S11nNode` and `Serializable` types. Anyone familiar with the STL's `char_traits<>` type will find the `s11n`-related traits types similar.

6.1 `s11n::node_traits<NodeType>`

Header file: `traits.hpp`

`node_traits` encapsulates the API of a given S-Node type. Using this approach it is possible to add new S-Node types to the framework without requiring clients to directly know about their concrete types. All that is required is a specialization of `node_traits` to act as the middleman between `s11n` and specific node types.

The complete API is documented in the `node_traits` API documentation.

Note that it is considered "poor form" to directly use the API of a given Node type in client code - use the traits type when possible.

The default `node_traits` implementation works with `s11n::s11n_node`. Using `node_traits` to manipulate these objects will ensure that client code can be used with either potential future node types.

It might be interesting to note that `s11n` has been used successfully with at least three node types, so the swapping-out-node-type idea has shown to be more than a theoretical feature.

6.2 `s11n::s11n_traits<SerializableType>`

Header file: `traits.hpp`

`s11n_traits` encapsulates the following information about a `Serializable` Type...

- **Serialization Functor** (typedef `serialize_functor`) - a functor type responsible for handling calls to `serialize()` on behalf of `SerializableType`.
- **Deserialization Functor** (typedef `deserialize_functor`) - a functor type responsible for handling calls to `deserialize()` on behalf of `SerializableType`. This is normally the same type as the **Serialization Functor**, but sometimes it may be necessary or desirable to implement different functors for each operation.

- Factory Type (typedef `factory_type`) - a functor which is responsible for creating new instances of the type (polymorphically, if required). This allows clients to easily install their own factories for a given class hierarchy, as opposed to being forced to use the default ones used by s11n.
- Cleanup Functor (typedef `cleanup_functor`) - added in 1.1.3 to allow some algorithms to make stronger guarantees in the face of exceptions. This functor is responsible for deallocating any otherwise unmanaged memory which might belong to a given type. This is used, e.g. to safely clean up containers containing pointers even when the pointers are nested in sub-containers.
- A single static function, `class_name(const serializable_type * HINT)`, added in version 1.1.0, allows algorithms to query Serializables for their class names in a more coherent way than in previous s11n versions (but with essentially the same effect and limitations *vis-a-vis* polymorphism).
- A “cleanup functor”, added in 1.1.3, which is conceptually similar to a destructor. This type is used to ensure proper cleanup up memory dynamically allocated during deserialization when exceptions are thrown (i.e., when s11n cannot give the allocated object(s) back to the user).

The interface and its conventions are documented fully in the `s11n_traits` API documentation.

Note that this type has no data members. That said, a specific traits specialization is free to expand the type. For example, it may contain the implementation for the de/serialization operators and typedef *itself* to be the `de/serialize_functor` types (yes, this has been done before and is perfectly kosher).

The original intention of `s11n_traits` was to replace SAM (section 17). As it turns out, SAM's (T*)-to-(T&) translation is fairly tricky to introduce via traits without an undue amount of extra code (potentially client-side). Since SAM does this in only a few lines of code, as is zero-maintenance (since early- or mid-2004 year), the pointer/reference translation support will stay in SAM. SAM is, however, implemented in terms of `s11n_traits`. That actually ends up giving us another layer we can hook in to, anyway, which gives us a bit more flexibility in swapping out components via specialization.

6.2.1 `cleanup_functor`

See also sections 19 and 16, which are closely related to this material.

This `s11n_traits`-specified type was added in 1.1.3 after realizing that this category of solution is the only way for the core library to avoid memory leaks in some particular cases involving failed deserialization.

In very specific terms, the job of the `cleanup_functor` is to deallocate resources which were dynamically allocated during deserialization. It is not intended to provide a general cleanup solution, only that necessary to free up memory allocated during deserialization transaction.

In short, this type is used to clean up factory-allocated objects if a deserialization involving those objects (directly or downstream) fails.

The cleanup functor is not normally directly used from client code unless the client has special needs in deserialization algorithms which require specific clean up in the face of failure. Even then, `s11n::cleanup_serializable()` is intended to act as a front-end to the cleanup functor.

Because failed deserialization normally leaves an object in an undefined state, we cannot simply delete such factory-allocated objects at will. The catch is, we don't know they're type, which means we might delete a `map<int, SomeT*>`, in which case a `delete` on the container would result in a leak of the `SomeT` members. Many of the major s11n algorithms are ignorant of pointeriness, and therefore don't even know if they're working with heap-allocated memory or not. They need a solution which can be used for heap- or stack-allocated objects using the same syntax, and so the `cleanup_functor` was developed.

For most client-side classes, those which manage their own memory (i.e., delete owned pointers at destruction), deleting the object on a failed deserialization is not a problem because it cleans its resources when it destructs. Deleting containers of *unmanaged* pointers is a severe problem, however.

There is a particular case for deserialization where the library cannot pass a newly-created object back to the caller (i.e., deser fails and the lib has an object it created). In that case, the library is forced to choose from three equally appalling choices:

1. Give back the object which failed deserialization. This option is not possible if an exception is thrown by the deser op, and in any case has no way of telling the caller that the object is in an undefined state. To the caller, it would seem as if all went well.
2. Don't delete the object, but give the user back null (meaning error), admitting a blatant leak.
3. Delete the object, admitting a leak only if the object contains unmanaged pointers.

Neither solution is satisfactory, but earlier versions of s11n had some failure cases which would take the third route (because the implications weren't recognized). Thank goodness deserialization failure at that level is so rare :/.

The `cleanup_functor` is expected to install rules for handling similar cases, such that on a deserialization failure we can internally call:

```
s11n_traits<SerializableT>::cleanup_functor() ( failedobject );
```

Assuming that functor does the right thing, that will clean up recursively on any contained elements, and any heap-allocated objects will be deleted. This does not happen all by itself - it requires conforming functors to be installed for each participating type. These are installed as part of the registration process, but special types will need some custom handling to install a proper cleaner-upper. Again, for PODs and classes which delete their member pointers at destruction, this is not an issue.

See the class template `s11n::default_cleanup_functor` for the API and required interface for specializations. Clients are not required to use that class, but it is the default implementation, and clients installing their own `s11n_traits` specializations must ensure that their cleanup functors behave as expected. You can find the various specializations installed for maps, pairs, and lists by grepping `proxy/*.hpp` for `default_cleanup_functor`. These might be useful starting points in writing your own, should you need to.

*Trivia: i delayed implementing `cleanup_functor` for some weeks because i was concerned about the build-time overhead the new required types would add (that's a sore point for me). On a small test i did using six binaries and two DLLs, the entire build time was only increased by about two seconds. The original prototype work for the approach was done almost a year before it was tried out here, but the larger implications of adding it never hit me until i actually started adding it (after finally realising that `(T * deserialize<T>(node))` was inherently leaky on failure of containers of pointers).*

6.3 `type_traits<T>`

Header file: `type_traits.hpp`

Version 1.1.2 introduced `type_traits<T>`, which is intended to be used by various algorithms to do things like stripping pointer and const qualifiers from types, and making compile-time decisions based off of such information. These types do not store any state and are not directly related to serialization other than as a utility to simplify some serialization code. There is nothing particularly special about this implementation - it is roughly similar to type traits found in many libraries.

7 Five-minute intro: PODs and STL containers

s11n's bread and butter is serializing PODs and STL containers. This short demonstration shows you everything you need to know to serialize most of them.

This whole section assumes the following typedefs, defining some client-side types we want to serialize:

```
typedef std::list<std::string> List;  
typedef std::map<int,List> ListMap;
```

And assumes we have some objects of those types:

```
List mylist;  
ListMap mymap;
```

It is irrelevant whether they are class members, globals, or whatever. As long as our code can access them, we don't care what scope they live in.

Reminder: the s11n source tree comes with many ready-to-run examples demonstrating a variety of common use cases: `src/client/sample/*.?pp`.

7.1 `#include ...`

First we need to include the core framework:

```
#include <s11n.net/s11n/s11n-lite.hpp>
```

Next we need to include a "proxy header" for each type we will de/serialize. These headers "promote" our types to Serializables (also called "registering" them).

Assuming the above-mentioned typedefs, we will need the following headers:


```
#include <s11n.net/s11n/proxy/std/list.hpp>
#include <s11n.net/s11n/proxy/std/map.hpp>
#include <s11n.net/s11n/proxy/pod/int.hpp>
#include <s11n.net/s11n/proxy/pod/string.hpp>
```

(Notice how the file names match the names of the type we want to serialize. This is a common convention.)

This *normally* equate to one header per *type* we want to serialize. Those last two (`pod/...`) headers aren't necessary in *some* cases, but we're going for the least-effort approach here, and the other approaches require knowing more about s11n than this intro assumes you currently know. Put briefly, we need those headers because the types *contained within a serializable container* must normally also be full-fledged Serializables, and we do this by including headers which install code to promote those PODs to Serializables. This is also why we can serialize `ListMap`, containing object of type `List`, as `List` is promoted to a `Serializable` via the inclusion of `list.hpp`. `ListMap` itself is promoted via `map.hpp`. The order of the includes is insignificant as long as all are included by the time we actually try to de/serialize objects of those type.

Trivia: by "promotion" to a Serializable, we mean taking a type which is not inherently Serializable and installing a proxy which acts on its behalf to provide Serializable behaviour. This allows us to non-intrusively add serialization features to many 3rd-party or built-in types, like the standard containers and built-in numeric types.

7.2 Saving

To save our objects, each one to its own file, is trivial:

```
s11nlite::save( mylist, 'list.s11n' );
s11nlite::save( mymap, 'map.s11n' );
```

Saving two disparate objects together inside *one* file requires a small bit more effort:

```
s11nlite::node_type node;
s11nlite::serialize_subnode( node, 'list', mylist );
s11nlite::serialize_subnode( node, 'map', mymap );
s11nlite::save( node, 'mystuff.s11n' );
```

7.3 Loading

Now let's load our objects:

```
List * l = s11nlite::load_serializable<List>( 'list.s11n' );
ListMap * m = s11nlite::load_serializable<ListMap>( 'map.s11n' );
```

If the loading fails, the pointers will be null or an exception may be thrown.

We can also deserialize directly from the file directly into an existing `List` or `ListMap` object:

```
std::auto_ptr<s11nlite::node_type> node( s11nlite::load_node( 'map.s11n' ) );
s11nlite::deserialize<ListMap>( *node, mymap );
```

Trivia: The explicit `<ListMap>` qualification on that `deserialize()` call is not necessary for monomorphic types, but it's a good habit to be in because it's often necessary to ensure proper s11n type lookup for polymorphs.

If we had saved both objects to one file, as shown above, we could load them with the following:

```
std::auto_ptr<s11nlite::node_type> node(
    s11nlite::load_node( 'mystuff.s11n' ) );
s11nlite::deserialize_subnode( *node, 'list', mylist );
s11nlite::deserialize_subnode( *node, 'map', mymap );
```

Notice how this time i left off the template type qualifiers. Again, this is fine for monomorphic types, but when writing generic de/serialization algorithms you should be in the habit of being explicit about the types.

7.4 Now the *really* easy way: `micro_api<>`

The obligatory header file:

```
#include <s11n.net/s11n/micro_api.hpp>
```

Create a micro:

```
s11nlite::micro_api<ListMap> mic;
```

Save/load:

```
mic.save( mymap, 'map.s11n' );  
ListMap * loaded = mic.load( 'map.s11n' );
```

Those are overloaded to take `iostream` objects.

If you don't need a file, don't bother with one. Instead, save it to a string buffer, which you can then save to a file, over a network, or to a copy/paste buffer:

```
mic.buffer( mymap );  
ListMap * loaded = mic.load_buffer();
```

The main advantage to the `micro_api` class is the elimination of all other template parameters involved with de/serialization. Another advantage is that the client code never needs to know about the "node type", which *very* is prevalent in the `s11n[lite]` APIs. The main limitation, however, is that each instance of `micro_api` is tied to a single *Serializable [Base] Type*, so we cannot use the same `micro_api` instance for both `mylist` and `mymap`. For many purposes, however, `micro` is the absolutely simplest way to save/load Serializables.

8 How to turn JoeAverageClass into a Serializable...

"... doing something about a problem which you do not understand is like trying to clear away the darkness by thrusting it aside with your hands."

Alan W. Watts, The Book (on the Taboo Against Knowing Who You Are)

Before we start: the `s11n` source tree and web site have a number of examples for using the library. You may want to check one of those places if this section does not help you.

In short, creating a `Serializable` is normally made up of these simple steps:

1. Create the class, implementing a pair of de/serialize methods with the signatures expected by `s11n`. The de/serialize operators may be defined in a separate (proxy) class in many common cases.
2. Tell `s11n` that your class exists, via registering it - see section 12.

If you are proxying a well-understood data structure for which a functor already exists to de/serialize it, step one disappears! An example would be proxying a `std::list<int>` or `std::list<Serializable*>` - those are both handleable by the `s11n::list::list_serializable_proxy` class, provided that the contained types are `Serializables`. For a list of some useful proxy functors see section 13. In the case of proxying standard containers, include the appropriate registration header file:

```
<s11n.net/s11n/proxy/std/list.hpp> // std::list  
<s11n.net/s11n/proxy/std/map.hpp> // std::map  
...
```

If the container contains types which must be proxied, those headers must also be included. For example, proxying a `map<int, string>` requires the following includes:

```
<s11n.net/s11n/proxy/std/map.hpp>  
<s11n.net/s11n/proxy/pod/int.hpp>  
<s11n.net/s11n/proxy/pod/string.hpp>
```

or an equivalent (there are other ways to do this). After that, any `std::map` containing any combination of ints or strings can be serialized via the core `s11n` API, including `map<string, int>` or `map<int, map<int, string>>`, etc.

8.1 Create a Serializable class

As you probably know by now, a Serializable's interface is made up two de/serialize operators. Types with different interfaces can also be used - see the next section. This library does not impose any inheritance requirements nor function naming conventions, but for this simple example we will take the approach of a serializable object hierarchy using the so-called *Default Serializable Interface*, made up of two overloaded `operator()`s.

Assume we've created these classes:

```
class MyType {
// serialize:
virtual bool operator()( s11n_lite::node_type & dest ) const;
// deserialize:
virtual bool operator()( const s11n_lite::node_type & src );
// ... our functions, etc.
};
class MySubType : public MyType {
// serialize:
virtual bool operator()( s11n_lite::node_type & dest ) const;
// deserialize:
virtual bool operator()( const s11n_lite::node_type & src );
// ... our functions, etc.
};
```

It is perfectly okay to make those operators member function *templates*, templated on the `NodeType`, but keep in mind that member function templates cannot be virtual. Implementing them as templates will make the serialization operators capable of accepting any Data Node type supported by s11n, which may have future maintenance benefits.

If a Serializable will not be proxied, as the ones shown above are not, we must register it as being a Serializable: see section 12 for how tell s11n about the class.

8.2 Specifying custom Serializable interfaces for InterfaceTypes

If `MyType` does not support the default interface, but has, for example:

```
[virtual] bool save()( data_node & dest ) const;
[virtual] bool load()( const data_node & src );
```

The library can still work with this. How to register the type as Serializable is described in section 12.

The same names may be used for both functions, as long as the constness is such that they can be properly told apart by the compiler.

8.3 Specifying Serializer Proxy functors

This is one of s11n's most powerful features. With this, any type can be made serializable *without editing the class*, provided its API is such that the desired data can be fetched and later restored. Almost all modern objects (those worth serializing) are designed this way, so this is practically never an issue.

Continuing the example from the previous section, if `MyType` cannot be made Serializable - if you can't, or don't want to, edit the code - then s11n can use a functor to handle de/serialize calls.

First we create a proxy, which is simply a struct or class with this interface:

Serialize:

```
bool operator()( DataNodeType & dest, const SerializableType & src ) const;
```

Deserialize:

```
bool operator()( const DataNodeType & src, SerializableType & dest ) const;
```

Notes about the operators:

- Yes, both functions "should probably" be `const` in this case, for the widest functor reusability, but if C++ will let you get away with non-`const` operators in your contexts then s11n will accept them.
- The operators may be templates and/or the functor may be a template. As long as C++'s type resolution can figure out what to do, it's legal.

- There are rare cases where calls can be ambiguously for this interface, so two functors - one each for de/serialization - may be necessary. (Trivia: in practice this has only once been necessary, and was probably caused by my mis-use of a non-const object.)

We must then register the proxy, as explained in section 12.6. For MyType and its subclass, shown above, the registration would look like this:

```
#define S11N_TYPE MyType
#define S11N_TYPE_NAME 'MyType'
// #define S11N_ABSTRACT_BASE // Only if MyType is abstract
#include <s11n.net/s11n/reg_s11n_traits.hpp>
#define S11N_TYPE MySubType
#define S11N_TYPE_NAME 'MySubType'
#define S11N_BASE_TYPE MyType
#include <s11n.net/s11n/reg_s11n_traits.hpp>
```

It may be interesting to know...

- There can be only one de/serialization handler for any given type, so you may not register both a base and a proxy as being the handler for a given type, nor may you register two proxies as being the proxy for a single Interface Type. Internally chaining calls within proxies can be used to get around the one-proxy limitation.
- Proxies may not normally save/load private data of the being-proxied type. In practice is is rarely an issue, as most modern libraries provide adequate accessors for their data. Classes designed such that they only possible way to store/restore their state is from internally should probably be redesigned to be more friendly. As a base-line comparison: every STL data structure which has been tried with this library has the necessary API to support proxying, with the exception of those with unusual traversal rules, like `queue` and `stack` (those two could be done, but would require an extra copy to be made, since we may not modify the source object during serialization).
- A proxy class does not need to register with a classloader. It may be registered - there is no harm in doing so, but there is never a need to³⁰. InterfaceType, on the other hand, must *always* be registered with the classloader.
- Proxies have a fixed interface - the function names and signatures may not be changed or marshaled (as Serializable interfaces can), for the simple reason that the proxies *are the ones doing the marshaling*.
- In theory it may sometimes be necessary, due to `const-vs-non-const` ambiguity, to split a de/serialization functor into two functors. In practice it's happened once, ever, back in s11n 0.7.x.
- Proxies can potentially chain calls to each other together, which allows some interesting possibilities and very flexible control over de/serialization without touching your classes. e.g. a data versioning system could be implemented as a proxy which introduces or verifies a version property and then passes on the call to the local Serializable interface of the object.
- Client code can, e.g. use a macro to define which proxy will be used for a given type (or group of types), allowing them to switch freely between serialization implementations on a per-type basis. This is how all of the "standard" proxies are implemented.

i have a feeling there are a wide range of as-yes-undiscovered tricks for serialization proxies. s11n early-adopter Gary Boone calls this feature "s11n's most powerful," and i can't help but agree with him.

9 How to turn JoeNonAverageClass into a Serializable...

"May your hands always be busy. May your feet always be swift. May you have a strong foundation when the winds of changes shift."

Bob Dylan

The techniques covered in the previous section work for most classes, but are not suitable for some others.

The following process works the same way for all types, as long as:

- It implements a serializable interface we can register with s11n.

or:

- A functor can be registered which will take over serilization for the type.

³⁰ Or, more correctly, if you understand the *highly unusual* (and purely theoretical) case that would warrant such registration, then you'll understand why we oversimplify here.

It is best shown with an example, where we proxy a client-supplied type:

```
#define S11N_TYPE MyType
#define S11N_TYPE_NAME "MyType"
// [de]serialization functor, only for proxied types:
#define S11N_SERIALIZE_FUNCTOR MyTypeSerializationProxy
// optional Derialization functor, defaults to S11N_SERIALIZE_FUNCTOR:
// #define S11N_DESERIALIZE_FUNCTOR MyTypeDeserializationProxy
#include <s11n.net/s11n/reg_s11n_traits.hpp>
```

You're done!

That's all that's necessary to take complete control over the internals of how s11n proxies a class.

This process must be repeated for each new type. The `S11N_XXX` macros are all unset after the registration header is included, so they may be immediately re-defined again in client code without having to undefine them first. Other proxy registration supermacros may implement whatever interface they like, with their own macro interfaces, allowing per-proxy-per-Serializable customization via macro toggles.

The registration process, on the surface, looks... well, *awkward*. Trust me, though: the benefits over of this simple approach macro- and code-generation-based solutions are tremendous, and have helped make some extremely tricky (or essentially impossible) cases much simpler to implement.

Note that when registering template types, you also need to register their *templated* types - they will be passed around just like other Serializables, so if s11n doesn't know about them you will get compile errors. And keep in mind that, e.g. `list<int>` and `list<int*>` *are different types*, and thus require different specializations. However, `list<int>` and `(list<int>*)` are equivalent for most of s11n's purposes.

9.1 JoeAverageClass<> class template

The s11n source tree contains a demonstration of this: `src/client/sample/templates.cpp`

Optionally, take a look at the standard proxies for the STL list/map containers, in the s11n source tree under `src/proxy/reg_{list,map}_specializations.hpp`. These files demonstrate the serialization proxying of class templates.

If you have a class template and a proxy prepared for it, you can register the template and its proxy with specialized supermacros dedicated to this purpose:

Template type with one templated parameter:

```
#define S11N_TEMPLATE_TYPE MyT
#define S11N_TEMPLATE_TYPE_NAME "MyT"
#define S11N_TEMPLATE_TYPE_PROXY MyT_s11n
#include <s11n.net/s11n/proxy/reg_s11n_traits_template1.hpp>
```

If MyT has two templated parameters:

```
#include <s11n.net/s11n/proxy/reg_s11n_traits_template2.hpp>
```

If you need to register a type with more than two parameters, you have at least two options:

1. Copy the `reg_s11n_traits_templateN.hpp` files into your tree and modify for more arguments. Please then send me a copy. :)
2. Wait until I personally need the feature, then the next s11n release will have it.

9.1.1 A cleanup functor

There is one additional concern when serializing class templates: if those types do not own pointers they contain then you must supply a "cleanup functor" so that the library knows how to deallocate your objects safely if it needs to as a result of an exception. To do this, simply provide a partial specialization of `s11n::default_cleanup_functor`, as briefly shown below and demonstrated in full in the sample source code.

Assuming MyT has two template parameters and is structured like a `std::pair`, we can implement a cleanup functor like this:

```
namespace s11n {
template <typename T1, typename T2>
```

```

struct default_cleanup_functor< MyT<T1,T2> > {
typedef MyT<T1,T2> cleaned_type;
void operator()( cleaned_type & c ) {
// example, assuming MyT is pair-like:
typedef typename ::s11n::type_traits<T1>::type _T1; // strip any pointer
typedef typename ::s11n::type_traits<T2>::type _T2; // ditto
::s11n::cleanup_serializable<_T1>( c.first );
::s11n::cleanup_serializable<_T2>( c.second );
}
};
}

```

Believe it or not, that works uniformly regardless of whether T1 and T2 are pointer types or not. We strip the pointer part so that if T1 or T2 are pointers, then the calls to `cleanup_serializable()` get *references to pointers*, which makes it capable of assigning those pointers to 0 after cleaning/deleting them.

Remember that the cleanup process is essentially a no-op for value/reference types, but deallocates pointers along the way. In the case of `MyT<int, MyT<long, string *>>`, cleaning up the outer-most MyT object will inherently climb down to clean up the (string*) part of the nested MyT. The same thing will happen for `MyT<A, B<C, M<K, V*>>>`, provided all of the nested types are Serializables with a proper cleanup functor installed. This ability is critical to guaranteeing no leaks in the face of exceptions.

The registration files for the standard containers also contain cleanup functor implementations which you can use as a basis for writing your own.

10 Doing things with Serializables

"...you aren't disappointed when using a DOS machine; you know what to expect and are pleasantly suprised if more happens."

Larry Anderson

Once you've got the Serializable "paperwork" out of the way, you're ready to implement the guts of your serialization operators. In s11n this is normally *extremely* simple. Some of the *many* possibilities are shown below.

In maintenance terms, the serialization operators are normally the only part of a Serializable which must be touched as a class changes. The "paperwork" parts do not change unless things like the class name or its parentage change [or you upgrade to a newer s11n which breaks old APIs or conventions].

Remember that when using Data Nodes, it is *strongly preferred* to use the `node_traits<NodeType>` interface, as opposed to the Node Type API directly, as explained in section 6.1. Client code may of course use typedefs to simplify usage of `node_traits`.

In the examples shown here we will assume the following typedef is in effect:

```

typedef s11n::node_traits<NodeType> NTR;

```

10.1 Setting "simple" properties

Any data which can be represented as a string key/value pair can be stored in a data node as a property:

```

NTR::set( node, 'my_property', my_value );

```

`set()` is a function template and accepts a string as a key and any *Streamable Type* as a value

There are cases involving ambiguity between ints/bools/chars which may require that the client explicitly specify the property's type as a template parameter:

```

NTR::set<int>( int, 'my_number', mynum );
NTR::set<bool>( node, 'my_number', mybool );

```

Each property within a node is unique: setting a property will overwrite any other property with the same name.

It must be re-iterated that `set()` **only** works when setting values which are *Streamable Types*. That is, types which support two complementary `ostream<<` and `istream>>` operators. To save Serializable children use the `serialize()` family of functions.

10.2 Getting property values

Getting properties from nodes is also very simple. In the abstract, it looks like:

```
T val = NTR::get( node, 'property_name', some_T_object );
```

e.g.

```
this->name( NTR::get( node, 'name', this->name() ) );
```

What this is saying is:

set this object's name to the value of the 'name' property of `node`. If 'name' is not set in `node`, or cannot be converted to a string via i/o streams, then use the current value of `this->name()`.

That sounds like a mouthful, but it's very simple: when calling `get()` you must specify a second parameter, which must be of the same type as the return result. This second parameter serves several purposes:

- A default value: a known-good (or known-bad!) value to use in case the supplied object could not be converted.
- An error value: The library cannot know what is and is not a valid value for such conversions, so the client may supply one here and compare it to what they expect. e.g. data versioning checks could be implemented this way.
- It tells `get()` what type of object it returns, without you having to specify `get<ReturnType>('mykey')`.

As with `set()`, `get()` is a family of overloaded/templated functions, and there are cases where, e.g. `int` and `bools` may cause ambiguity at compile time. See the `set()` documentat, above, for the proper workaround.

As with `set()`, `get()` **only** works with *Streamable Types*. To restore Serializable children, use the `deserialize()` family of functions.

You can also use `NTR::is_set(node, 'property')` to check for existence of a property.

10.2.1 Simple property error checking

Here's how one might implement simple error checking for properties:

```
int foo = NTR::get( node, 'meaning_of_life', -1 );
if( -1 == foo ) { ... error: we all know its really 42 ... }
std::string bar = NTR::get( node, 'name', std::string() );
if( bar.empty() ) { ... error ... }
if( ! NTR::is_set(node, 'important') ) { ... error ... }
```

Keep in mind that `s11n` cannot know what values are acceptable for a given property, thus it can make no assumptions about what values might be invalid or error values.

Theoretically, installing a Serializable Proxy for a type which does such checks and then passes the call on to the object's local Serializable Interface is one way to keep this type of code out of Serializable classes.

10.2.2 Saving custom Streamable Types

This is a no-brainer. Streamable Types are supported using the same `get/set` interface as all other "simple" properties. Assume we have a `Geometry` type which support i/o stream operators. In order to save it we must simply call:

```
NTR::set( node, 'geom', this->geometry() );
```

and to load it:

```
this->geometry( NTR::get( node, 'geom', this->geometry() ) );
```

or maybe:

```
this->geometry( NTR::get( node, 'geom', Geometry() ) );
```

10.3 Finding or adding child nodes to a node

Use the `s11n::find_child_by_name()` and `s11n::find_children_by_name()` functions to search for child nodes within a given node. Alternately, use `node_traits<NodeType>::children()` function to get the list of its children, and search for them using criteria of your choice.

Use `s11n::create_child()` to create a child and add it to a parent in one step. Alternately, add children using `node_traits<NodeType>::children(node).push_back()`.

10.4 Serializing Streamable Containers

Streamable Containers are, in this context, containers for which all stored types are *Streamable Types* (see 4.1). `s11n` can save, load, and convert such types with unprecedented ease.

Normally containers are stored as sub-nodes of a Serializable's data node, thus saving them looks like:

```
s11n::map::serialize_streamable_map( node, 'subnode_name', my_map );
```

To use this function directly on a target node, without an intervening subnode, use the two-argument version without the subnode name. Be warned that none of the `serialize_xxx()` functions are meant to be called repeatedly or collectively on the same data node container. That is, each one expects to have a "private" node in which to save its data, just as a full-fledged Serializable object's node would. Violating this may result in mangled content in your data nodes, or possibly an exception, depending on the algo (in 1.1.3+ most algos throw in this case).

Loading a map requires exactly two more characters of work:

```
s11n::map::deserialize_streamable_map( node, 'subnode_name', my_map );
```

(Can you guess which two characters changed? ;)

If you want to de/serialize a `std::list` or `std::vector` of Streamable Types, use the `de/serialize_streamable_list()` variants instead:

```
s11n::list::serialize_streamable_list( targetnode, 'subnodename',  
my_list );
```

Note that `s11n` does not store the *exact* type information for data serialized this way, which makes it possible to convert, e.g. a `std::list<int>` into a `std::vector<double*>`, via serialization. The wider implication is that any list- or map-like types can be served by these simple functions (all of them are implemented in 6-8 lines of code, not counting typedefs). We actually rely on C++'s strong typing to do the hardest parts of type determination, and we don't actually need the type name in some cases involving monomorphic Serializables. More specifically, whenever no classloading operation is required, the class name *ist uns egal*³¹.

Note that these functions only work when the contained types are Streamables. If they are not, use the `s11n::list::serialize_list()` and `s11n::map::serialize_map()` family of functions. Note that those functions also work for Streamable types as long as a proxy has been installed for those Streamables (see `proxy/pod/*.hpp` for examples).

10.4.1 Trick: "casting" list or map types

If you have lists or maps which are similar, but not exactly of the same types, `s11n` can act as a middleman to convert them for you. Assume we have the following maps:

```
map<int,int> imap;  
map<double,double> dmap;
```

We can convert `imap` to `dmap` like this:

```
data_node n;  
s11n::map::serialize_streamable_map( n, imap );  
s11n::map::deserialize_streamable_map( n, dmap );
```

In fact, that doesn't require that any of the involved types be registered Serializables, provided the algorithms' other requirements are met.

For Serializables we have a simpler option:

³¹ German for "frankly, my dear, we don't give a damn."

```
s11nlite::s11n_cast( imap, dmap );
```

This requires that proxies be in place for the maps as well as the contained types, `int` and `double`, which we can install with:

```
#include <s11n.net/s11n/proxy/std/map.hpp>
#include <s11n.net/s11n/proxy/pod/int.hpp>
#include <s11n.net/s11n/proxy/pod/double.hpp>
```

Doing the opposite conversion via `s11n_cast()` "should" also work, but would be a potentially bad idea because any post-decimal data of the `doubles` would be lost upon conversion to `int`. The compiler cannot warn you about loss of precision in such a case because the conversions happen via lexical casting.

Similar conversions will work, for example, for converting a `std::list` to a `std::vector`. For example:

```
#include <s11n.net/s11n/proxy/std/list.hpp>
#include <s11n.net/s11n/proxy/std/vector.hpp>
#include <s11n.net/s11n/proxy/pod/int.hpp>
...
list<int>  ilist;
vector<int*>  ivec;
// ... populate ilist ...
s11nlite::s11n_cast( ilist, ivec );
```

That's all there is to it. The library takes care of allocating the (`int*`) children of the vector. The client is responsible for deallocating them, just as one would when using any "normal" STL container of pointers. One simple way to deallocate them:

```
s11n::cleanup_serializable( ivec );
```

That works even if the vector contains containers which contain containers which themselves contain more containers of pointers.

10.5 De/serializing Serializable objects

In terms of the client interface, saving and restoring Serializable objects is slightly more complex than working with basic types (like PODs), primarily because we must deal with more type information.

10.5.1 Individual Serializable objects

The following C++ code will save any given Serializable object to a file:

```
s11nlite::save<MyType>( myobject, 'somefile.whatever' );
```

this will save it into a target `s11nlite::node_type` object:

```
s11nlite::serialize<MyType>( mynode, myobject );
```

The node could then be saved via an overloaded form of `save()`.

There are several ways to save a file, depending on what Serializer you want to use. `s11nlite` uses only one Serializer by default, so we'll skip that subject for now (tips: see `s11nlite::serializer_class()` for a way to override which Serializer it uses).

Loading an object is fairly straightforward. The simplest way is:

```
InterfaceType * obj =
    s11nlite::load_serializable<InterfaceType>( 'somefile.s11n' );
```

`InterfaceType` must be a type registered with the appropriate classloader (i.e., the `InterfaceType` classloader) and must of course be a Serializable type. To illustrate that first point more clearly, **the following are not correct**:

```
SubTypeOfInterfaceType * obj =
    s11nlite::load_serializable<InterfaceType>( 'somefile.s11n' );
```

Will not compile: there is no implicit conversion from `InterfaceType` to a subtype of that type.

```
InterfaceType * obj =
    s11nlite::load_serializable<SubTypeOfInterfaceType>('somefile.s11n');
```

Will compile but will not do what is expected, because it's trying to use a different classloader and API marshaller than `InterfaceType`.

It is critical that you use the base-most type which was registered with `s11n`, or you will almost certainly not get back an object from any deserialize-related function.

If you have a non-pointer type which must be populated from a file, it can be deserialized by getting an intermediary data node, by using something like the following:

```
s11nlite::node_type * n = s11nlite::load_node( 'somefile.s11n' );
```

or:

```
const s11nlite::node_type * n = s11n::find_child_by_name( parent_node,
    'subnode_name' );
```

Then, assuming you got a node:

```
bool worked = s11nlite::deserialize( *n, myobject );
delete( n );
// ^^^^ NOT if you got it from another node! It belongs to the parent node!
```

Note, however, that if the deserialize operation fails then `myobject` might be in an undefined or unusable state. In practice this is *extremely rare*, but it may happen, and client code may need to be able to deal with this possibility.

10.5.2 Containers of Serializables

This subsection exists only to avoid someone asking, "how do I serialize a `list<T>` or `list<T*>`?"

Here you go:

```
#include <s11n.net/s11n/proxy/listish.hpp> // list-related algos
#include <s11n.net/s11n/proxy/std/list.hpp> // proxy registration
...
s11n::serialize( target_node, src_list );
...
s11n::deserialize( src_node, tgt_list );
// or:
ListType * tgt_list = s11n::deserialize<ListType>( src_node );
```

The same goes for maps, except that you should include `mapish.hpp` and `std/map.hpp`. Note that "list" algorithms actually work with `std::list`, `vector`, `set` and `multiset`, but that proxies for each general list type must be installed separately, by including one of `std/{list,set,vector,...}.hpp`. The map algorithms work for `std::map` and `multimap` and are proxied via the headers `std/{multimap,map}.hpp`.

So what is different from the above code and de/serialization of any other Serializable type? Nothing. That's part of what makes `s11n` so easy to use - clients only really need to remember a small handful of functions.

10.5.3 "Brute force" deserialization

Any data node can be de/serialized into any given Serializable, provided the Serializable supports a deserialize operator for that node type. The main implication of this is that clients may force-feed any given node into any object, regardless of the meta-data type of the data node (i.e., its `class_name()`) and the Serializable's type. This feature can be used and abused in a number of ways, and one of the most common uses is to deserialize non-pointer Serializables:

```
if( const data_node * ch = s11n::find_child_by_name( srcnode, 'fred' ) ) {
    if( ! s11nlite::deserialize<MyType>( *ch, myobject ) ) {
        ... error ...
    }
}
```

The notable down-side of brute-force deserialization, however, is this: if the deserialize operation fails then `myobject` *may* be in an undefined state, depending on the algorithm used to deserialize it. Handling of this is

(a) very client-specific, and (b) in practice it is very rare for a deserialization to fail at this level. Brute force deserialization specifically opens up the possibility of feeding any data to any deserialization algorithm, which of course means that for correct results you must use matching data and algorithms.

11 Walk-throughs: implementing Serializable classes

This section contains some example of implementing real-world-style Serializables. It is expected that this section will grow as exceptionally illustrative samples are developed or submitted to the project.

There are several complete, documented examples in the source tree under `src/client/...`, and the s11n web site has several. Both sources go well beyond what is presented here.

11.1 Sample #1: Read this before trying to code a Serializable!

Here we show the code necessary to save an imaginary client-side Serializable class, `MyType`.

The code presented here could be implemented either in a Serializable itself or a in a proxy, as appropriate. The code is the same, either way.

In this example we are not going to proxy any classes, but instead we will use various algorithms to store them. The end effect is identical, though the internals of each differ slightly.

11.1.1 The data

Let's assume that `MyType` has this rather ugly mix of internal data we would like to save:

```
std::map<int,std::string> istrmap;
std::map<double,std::string> dstrmap;
std::list<std::string> slist;
std::list<MyType *> childs;
size_t m_id;
```

Looks bad, doesn't it? Don't worry - this is a trivial case for s11n.

11.1.2 The #includes

We will need to include the following headers for our particular case:

```
#include <s11n.net/s11n/s11n-lite.hpp>
#include <s11n.net/s11n/proxy/std/list.hpp> // list proxy
#include <s11n.net/s11n/proxy/std/map.hpp> // map proxy
#include <s11n.net/s11n/proxy/pod/int.hpp> // see below
#include <s11n.net/s11n/proxy/pod/double.hpp> // see below
#include <s11n.net/s11n/proxy/pod/string.hpp> // see below
```

The `pod/xxx.hpp` headers promote the given PODs to first-class Serializables. This is not necessary, nor desirable, for all cases, but simplifies this example.

11.1.3 The serialize operator

Saving member data normally requires one line of code per member, as shown here:

```
bool operator()( s11n-lite::node_type & node ) const
{
    typedef s11n-lite::node_traits TR;
    TR::class_name( node, "MyType" ); // critical, but see below!
    TR::set( node, "id", m_id );
    using namespace s11n-lite;
    serialize_subnode( node, "string_list", slist );
    serialize_subnode( node, "children", childs );
    serialize_subnode( node, "int_to_str_map", istrmap );
    serialize_subnode( node, "dbl_to_str_map", dstrmap );
    return true;
}
```

The class name for a registered monomorphic Serializable types can be fetched by calling

`::classname<T>()`. In fact, SAM (section 17) does this for you, and the `class_name()` call can technically be left out for monomorphic types. It is probably a good idea to go ahead and include it, for the sake of clarity and pedantic correctness.

If we had not promoted our PODs to first-class serializables, using `pod/xxx.hpp`, we could still serialize our data, but would then need create registrations to map them to specific proxies or call the desired algorithms ourselves. Both are desirable under particular circumstances. A sample of how that might be done:

```
s11n::list::serialize_streamable_list( node, "string_list", slist );
s11n::map::serialize_streamable_map( node, "int_to_str_map", istrmap );
```

Those algorithms produce much more compact output than the default proxies, but are only useful when all types contained in the container are `i/ostreamable`.

11.1.4 The deserialize operator

The deserialize implementation is almost a mirror-image of the serialize implementation, plus a couple lines of client-dependent administrative code (not always necessary, as explained below):

```
bool operator()( const s11n_lite::node_type & node )
{
    /////////////////////////////////////////////////// avoid duplicate entries in our lists:
    istrmap.clear();
    dstrmap.clear();
    slist.clear();
    s11n::cleanup_serializable( this->childs );
    /////////////////////////////////////////////////// now get our data:
    typedef s11n_lite::node_traits TR;
    this->m_id = TR::get( node, "id", m_id );
    using namespace s11n_lite;
    deserialize_subnode( node, "string_list", slist );
    deserialize_subnode( node, "children", childs );
    deserialize_subnode( node, "int_to_str_map", istrmap );
    deserialize_subnode( node, "dbl_to_str_map", dstrmap );
    // ^^ If we previously used serialize_streamable_xxx() we would
    // need to use deserialize_streamable_xxx() to retrieve the data.
    return true;
}
```

A note about cleaning up *before* deserialization:

In practice these checks are normally not necessary. `s11n` never, in the normal line of duty, directly calls the deserialize operator more than one time for any given `Serializable`: it calls the operator one time directly after instantiating the object. It is conceivable, however, that client code will initiate a second (or subsequent) deserialize for a live object, in which case we need to avoid the possibility of appending to our current properties/children, and in the above example we avoid that problem by clearing out all children and lists/maps first. In practice such cases tend to only happen in test/debug code, not in real client use cases. The possibility of multiple-deserialization *is* there, and it is potentially ugly, so it is prudent to add the extra few lines of code necessary to make sure deserialization starts in a clean environment.

11.1.5 Serializable/proxy registration

The interface must now be registered with `s11n`, so that it knows how to intercept requests on that type's behalf: for full details see section [12](#), and for a quick example see [9](#).

11.1.6 Done! Your object is now a Serializable Type!

That's all there is to it. Now `MyType` will work with any `s11n` API which work with `Serializables`. For example:

```
s11n_lite::save( myobject, std::cout );
```

will dump our `MyObject` to `cout` via `s11n` serialization. This will load it from a file:

```
MyType * obj = s11n_lite::load_serializable<MyType>( 'filename.s11n' );
```

(Keep in mind that the object you get back might actually be some ancestor of `MyType` - this operation is polymorphic if `MyType` is.)

Now that wasn't so tough, was it?

A very significant property of MyType is this:

MyType is now inherently serializable by any code which uses s11nlite, regardless of the code's local Serialization API! s11n takes care of the API translation between the various local APIs.

Weird, eh? Let's take a moment to day-dream:

Consider for a moment the *outrageous* possibility that 746 C++ developers worldwide implement s11n-compatible Serializable support for their objects. Aside from having a convenient serialization library at their disposal (i mean, *obviously* ;) , those 746 developers now have *100% transparent* access to each others' serialization capabilities, without having to know anything but the other libraries' base-most types.

Now consider for a moment the implications of *your* classes being in that equation...

Let us toke on that thought for a moment, absorbing the implications.

Well, *i* think it's pretty cool, anyway.

11.2 Gary's code

One of s11n's early-adopters, Gary Boone, contacted me in early 2004 about how to go about adding s11n support to his project. For starters, he had a simple structure (described below). On the surface, the problem appears to be non-trivial, but this is only when viewing the code through the lense of traditional C++ techniques...

Let us repeat the s11n mantra (well, one of several³²):

s11n is here to Save Our Data, man!

The type of problem Gary is trying to solve here is s11n's *bread and butter*, as his solution will show us in a few moments.

After getting over the initial learning hurdles - admittedly, s11n's abstractness can be a significant hinderness in understanding it - he got it running and sent me an email, which i've reproduced below with his permission.

i must say, it gives me *great pleasure* to post Gary's text here. Through his mails i have witnessed the dawning of his excitement as he comes to understanding the general utility of s11n, and that is one of the greatest rewards i, as s11n's author, can possibly get. Reading his mails certainly made me feel good, anyway :).

Gary's email address has been removed from these pages at his request. If, after reading his examples, you're intested in contacting Gary, please send me a mail saying so and i will happily forward it on to him.

The code below has been updated from Gary's original to accomodate changes in the core library, but it is essentially the same as his original post.

In some places i have added descriptive or background information, marked like so:

[editorial:]

11.2.1 Gary's Revelation

[From: Gary Boone, 12 March 2004]

...

Attached is my solution ('map_of_structs.*'). Basically, I followed your suggestion of writing the vector elements as node children using a for_each & functor.

...

I like the idea of not having to change **any** of my objects, but instead use functors to tell s11n how to serialize them.

...

Dude, it works!! That's amazing! That's huge, allowing you to code serialization into your projects without even touching other people's code in distributed projects. It means you can experiment with the library without having to hack/unhack your primary codebase.

32 Trivia note: The banner label on the s11n web site rotates through s11n's list of official mantra, and new mantra are added as they ar discovered. Submit your s11n mantra or clever quip and it will show up on the s11n web site. :)

Stephan, you **have** to make this clearer in the docs! It should be example #1:

[editorial: i feel compelled to increase the font size of that last part by a few points, because i had the distinct impression, while reading it, that Gary was overflowing with amazement at this realization, just as i first did when the implications of the architecture started to trickle in. :) That said, the *full* implications and limits of the architecture not yet fully understood, and probably won't be in the foreseeable future - i honestly believe it to be *that flexible*³³.]

...

One of the most exciting aspects of s11n is that you may not have to change **any** of your objects to use it! For example, suppose you had a struct:

```
struct elem_t {
    int index;
    double value;
    elem_t(void) : index(-1), value(0.0) {}
    elem_t(int i, double v) : index(i), value(v) {}
};
```

You can serialize it without touching it! Just add this proxy functor so s11n knows how to serialize and deserialize it:

```
// Define a functor for serialization/deserialization
// of elem_t structs:
struct elem_t_s11n34 {
    // note: no inheritance requirements, but
    // polymorphism is permitted.
    /*****
    // a so-called 'serialization operator':
    // This operator stores src's state into the dest data container.
    // Note that the SOURCE Serializable is const, while the TARGET
    // data node object is not.
    *****/

    template <typename NodeType>
    bool operator()( NodeType &dest, const elem_t &src ) const35 {
        typedef s11n::node_traits<NodeType> TR;
        TR::class_name( dest, "elem_t");
        TR::set( dest, "i", src.index);
        TR::set( dest, "v", src.value);
        return true;
    }

    /*****
    // a 'deserialization operator':
```

33 That text was written some time in the 0.7 or 0.8 cycle, early 2004 (today == 24 Sept 2005). i still believe that (a) the full limits and implications of the library are not yet fully understood and (b) it really is that flexible. :)

34 Gary is credited with coming up with the MyType_s11n naming scheme, and it now appears regularly in other s11n client trees.

35 Whether or not a functor has const or non-const operator(s) is largely a matter of what the functor is used for. The constness of the *arguments* is *set* - it may not deviate from that shown here. The constness of the operator itself is not defined by s11n conventions.

```

// This operator restores dest's state from
// the src data container.
// Note that the SOURCE node is const, while
// the TARGET Serializable object is not.
*****/

template <typename NodeType>
bool operator() ( const NodeType & src, elem_t & dest ) const {
    typedef s11n::node_traits<NodeType> TR;
    dest.index = TR::get( src, "i", -1);
    dest.value = TR::get( src, "v", 0.0);
    return true;
}
};

```

[editorial: while the similar-signatured overloads of `operator()` may seem confusing or annoying at first, with only a little practice they will become second nature, and the symmetry this approach adds to the API improves its overall ease-of-use. Note the bold text in their descriptions, above, form simple pneumonics to remember which operator does what.

The constness of the arguments ensures that they cannot normally (i.e., via standard s11n operations) be called ambiguously. That said, i have seen *one* case of a proxy *functor* (not Serializable) for which const/non-const-ambiguity was a problem, which is why proxies *may* optionally be implemented in terms of two objects: one SerializeFunctor and a corresponding DeserializeFunctor, each of which must implement their corresponding halves of the de/serialize equation. Often it is very useful to first implement de/serialize *algorithms* (i.e. *as functions*) and then later supply the 8-line wrapper *functor* class which forwards the calls to the algorithms. Several internal proxies do exactly this, and it gives client code two different ways of doing the same thing, at the cost of an extra couple minutes of coding the proxy wrapper around an existing algorithm. As a general rule, algorithms are slightly easier to test than proxies early on in development, as they are missing one level of indirection which proxies logically bring along.

Back to you, Gary...

The final step is to tell s11n about the association between the proxy and its delegatee:

```

#define S11N_TYPE elem_t
#define S11N_TYPE_NAME 'elem_t'
#define S11N_SERIALIZE_FUNCTOR elem_t_s11n
#include <s11n.net/s11n/reg_s11n_traits.hpp>

```

[editorial: After this registration, `elem_t_s11n` is now *the* official delegate for *all* de/serialize operations involving `elem_t`. Any time a de/serialize operation involves an `elem_t` or (`elem_t *`) s11n will direct the call to `elem_t_s11n`. The only way for a client to bypass this proxying is to do the most *dispicable*, *unthinkable* act in all of libs11n: passing the node to the Serializable directly using the Serializable's API! See section 5.4 for an explanation of why taking such an action is considered *Poor Form!*]

You're done. Now you can serialize it as easily as:

```

elem_t e(2, 34.5);
s11nlite::save(e, std::cout);

```

Deserializing from a file or stream is just as straightforward:

```

elem_t * e = s11nlite::load_serializable<elem_t>( "somefile.elem" );
or:
s11nlite::data_node * node = s11nlite::load_node( "somefile.elem" );
elem_t e;
bool worked = s11nlite::deserialize( *node, e );

```

```
delete node;
```

[editorial: that last example basically "cannot fail" unless `elem_t`'s `deserialize` implementation *wants* it to, e.g. if it gets out-of-range/missing data and decides to complain by returning false. What might cause *missing* data in a node? That's exactly what would effectively happen if you "brute-force" a node populated from a non-`elem_t` source into `elem_t`. Consider: the node will probably *not* be laid out the same internally (different property names, for example), and if it *is* laid out the same, there are still no guarantees such an operation is *semantically* valid for `elem_t`. Obviously, handling such cases is 100% client-specific, and must be analyzed on a case-by-case basis. In practice this problem is mainly theoretical/academic in nature. Consider: frameworks understand their own data models, and don't go passing around invalid data to each other. s11n's strict classloading scheme means it cannot inherently do such things, so that type of "use and abuse" necessarily comes from client-side code. Again: *this never happens*. Jesus, i'm so pedantic sometimes...]

...

[End Gary's mail]

Gary hit it right on the head. The above code is *exactly* in line with what s11n is designed to do, and his first go at a proxy was implemented exactly correctly. Kudos, Gary!

Note that with the various container proxies which ship with s11n, Gary's `elem_t` type can take part in container serialization, such as in a `map<string,elem_t>`

or `list<elem_t>`. There is no separate "serialize container of `elem_t`" operation, as the generic list/map algorithms inherently handle any and all `Serializable`s:

```
typedef std::map<std::string,elem_t> MapT;
MapT mymap;
... populate mymap ...
s11nlite::save( mymap, 'myfile.s11n' );
```

12 s11n registration & "supermacros" (IMPORTANT)

As of version 0.8.0, s11n uses a new class registration process, providing a single interface for registering any types, and handling all classloader registration.

Historically, macros have been used to handle registration, but these have a *huge* number of limitations. We now have a new process which, while a tad more verbose, is far, far superior in many ways (the only down-side being its verbosity). i like to call them...

12.1 "Supermacros"

s11n uses generic "supermacros" to register anything and everything. A supermacro is a header file which is written to work like a C++ macro, which essentially means that it is designed to be passed parameters and included, potentially repeatedly.

Use of a supermacro looks something like this:

```
#define MYARG1 'some string'
#define MYARG2 foo::AType
#include 'my_supermacro.hpp'
```

By convention, and for client convenience, the supermacro is responsible for unsetting any arguments it expects after it is done with them, so client code may repeatedly call the macro without `#undef`'ing them.

Sample:

```
#define S11N_TYPE MyType
#define S11N_TYPE_NAME "MyType"
#define S11N_SERIALIZE_FUNCTOR MyType_s11n
#include <s11n.net/s11n/reg_s11n_traits.hpp>

#define S11N_TYPE MyOtherType
#define S11N_TYPE_NAME "MyOtherType"
#define S11N_SERIALIZE_FUNCTOR MyOtherType_s11n
#include <s11n.net/s11n/reg_s11n_traits.hpp>
```

While the now-outmoded registration macros are (barely) suitable for many non-templates-based cases, supermacros allow some - er... *TONS* - of features which the simpler macros simply cannot come close to

providing. For example:

- A supermacro can handle almost any case, using a single - yet extendable - interface, and more complex variants can implement their own supermacro file.
- Supermacros can do arbitrary tasks, like classloader registration, freeing clients of this task.
- Arbitrary new sets of supermacros can be introduced at any time without impacting existing code, which means, for example, client code can use a `#define` to switch between interfaces by including different registration macros.
- ODR violations can be more easily eliminated (in theory, completely), as each supermacro is free to implement its internals however it wants. e.g. if it uses a custom classloader registration technique it cannot collide with those used by other registerers.
- As they are implemented in "real header code", they are completely immune to the limitations of macros, and simply *much* easier to maintain.
- This approach does ALL necessary registration, including classloader registration (could not be reliably done via the macro approach, due to ODR-violation possibilities).
- Supermacros can be arbitrarily large, whereas macros get very tedious to edit once they are longer than a few lines.
- They are *much, much* easier to debug when something doesn't compile: unlike conventional macros, we even get proper file names and line numbers (*yes!!!!*).

The adoption of the supermacro mechanic into s11n 0.8 opened up a huge number of possibilities which were simply not practical to do before, and implications are still not fully appreciated/understood.

12.2 General: Interface Types

All of s11n's activity is "keyed" to a type's Interface Type. This is used for a number of internal mechanisms, far too detailed to even properly *summarize* here. A `InterfaceType` represents the base-most type which a "registration tree" knows about. In client/API terms, this means that when using a hierarchy of types, the base-most `Serializable` type should be used for all templated `InterfaceType/SerializableType` parameters.

(See, it's difficult to describe!)

In most usage using `InterfaceTypes` as key is quite natural and normal, but one known case exists where they can be easily confused:

Assume we have this heirarchy:

```
AType <-[extended by]- BType <-- CType
```

In terms of matching `InterfaceType` to subtypes, for *most purposes*, that looks like this:

- BType's `InterfaceType` is AType
- CType's `InterfaceType` is **AType**

There are valid cases where registering both AType and BType as bases of CType are useful, but doing so in the same compilation unit will fail with the default registration process, with ODR collisions. The need to do this is rare (or non-existent, for most practical purposes), in any case, and requires a good understanding of how the classloader works. Doing it is very straightforward, but requires a bit of client-side effort.

12.3 Choosing class names when registering

s11n does not care what class names you use. We could call, e.g. `std::map<string, string>` "fred" and the end effect is the same. In fact, we could also call the pair type contained in that map "fred" - *without getting a collision* - because it uses a different classloader than the map (because they have different `InterfaceTypes`, as described in section 12.2).

The important thing is that we are consistent with class names. Once we change them, any older data will not be loadable via the classloader unless we explicitly alias the type names via the factory's aliasing API (see `s11n::cl::classloader_alias()`).

By convention, s11n uses a class' C++ name, stripped of spaces and any `const` and pointer parts. The "noise" parts are, it turns out, *irrelevant* for purposes of classloading and cause *completely unnecessary* maintenance in other parts of the code (including, potentially, client code). Thus, when s11n saves a `(std::string)` or a `(std::string*)` the type name s11n uses will be "std::string" (or even "string") for *both* of them, and the context of the de/serialization determines whether we need to dynamically allocate pointers or not. It is, of course, up to client code to deallocate any pointers created this way. For example, when deserializing a `list<string*>`, the client must free the list entries. (Tip: see `s11n::cleanup_serializable()` for a simple, generic way to accomplish this.)

12.4 Registering Interface Types supporting serialization operators

As of s11n 0.8, s11n "requires" so-called Default Serializables to be registered. In truth, they don't *have* to be for all cases, but for widest compatibility and ease of use, it is highly recommended. It is pretty painless, and must be done only one time per type:

```
#define S11N_TYPE ASerType
#define S11N_TYPE_NAME "ASerType"
#include <s11n.net/s11n/reg_s11n_traits.hpp>
```

The registration of a subtype of ASerType looks like:

```
#define S11N_BASE_TYPE ASerType
#define S11N_TYPE BSerType
#define S11N_TYPE_NAME "BSerType"
#include <s11n.net/s11n/reg_s11n_traits.hpp>
```

The S11N_XXX macros are #undef'ed by the registration code, so client code need not do so, and may register several classes in a row by simply re-defining them before including the supermacro code.

12.5 Registering types which implement a custom Serializable interface

If a class implements two serialization functions, but does not use operator() overloads, the process is simply a minor extension of the default case described in the previous section. We must do two things:

First, define a functor which, in its Serialization Operators, forwards the call to MyType's serialization interface. An example of such a functor:

```
struct MyType_s11n {
// note that the proxy class name is unimportant:
// Gary Boone came up with the XXX_s11n convention
// and i adopted it.
    template <typename NodeType>
    bool operator()( NodeType & dest, const MyType & src ) const {
        return src.local_serialize_function( node );
    }
    template <typename NodeType>
    bool operator()( const NodeType & dest, MyType & src ) const {
        return src.local_deserialize_function( node );
    }
};
```

Second, before including the registration supermacro as shown in the previous section, simply add one or both of these defines:

```
#define S11N_SERIALIZE_FUNCTOR MyType_s11n
#define S11N_DESERIALIZE_FUNCTOR MyType_s11n
// ^^^^ OPTIONAL: defaults to S11N_SERIALIZE_FUNCTOR
```

The second functor is only necessary if you define *separate functor classes* for de/serialization operations. In the vast majority of cases one proxy handles both de/serialize operations, so the second macro need not be set.

That's it - you're done telling s11n how to talk to your local serialization API. Now calls to s11n::de/serialize() will end up routing through the local_de/serialize_function() API.

12.6 Registering Serializable Proxies

In fact, there is no one single way to do this, because there are several pieces to a registration:

The important things are:

- *Proxied (not proxy)* type must be registered with appropriate classloader: monomorphs register with their own, as do Interface/Base-most Types, and subclasses register with their Interface Type's classloader.
- s11n_traits<ProxiedType>::class_name() should return the class name which s11n will use for the type. For monomorphs the library can figure this out on its own, but needs help with polymorphic type names. As a general rule, this cannot be achieved in this function for polymorphs,

which is why we say that serialize operators for polymorphs must always explicitly set their correct (sub)type name.

- An `s11n_traits<>` specialization installed (section 6.2).

After months of experimentation, `s11n` refines the process to simply calling the following supermacro:

```
#define S11N_TYPE ASerType
#define S11N_TYPE_NAME "ASerType"
#define S11N_SERIALIZE_FUNCTOR ASerType_s11n
// optional: #define S11N_DESERIALIZE_FUNCTOR ASerType_des11n
// DESERIALIZE defaults to the SERIALIZE functor, which works fine for most
cases.
#include <s11n.net/s11n/reg_s11n_traits.hpp>
```

Note that the names of the de/serialize functors shown here are arbitrary: you'll need to use the name(s) of *your* proxy type(s).

This is repeated for each proxy/type combination you wish to register. The macros used by `reg_s11n_traits.hpp` are temporary, and `#undef'd` when it is included.

There are other optional macros to define for that header: see `reg_s11n_traits.hpp` for full details.

If we extend `ASerType` with `BSerType`, B's will look like this:

```
#define S11N_BASE_TYPE ASerType
#define S11N_TYPE BSerType
#define S11N_TYPE_NAME "BSerType"
#include <s11n.net/s11n/reg_s11n_traits.hpp>
```

Without the need to specify the functor name - it is inherited from the type set in `S11N_BASE_TYPE`.

12.7 Registering a proxy for a class template

The library has had, for quite some time, a few supermacros to help out with the creation of proxy code for class templates. As of version 1.2.5, this includes supermacros with up to 4 template arguments and templates which use concrete types, as opposed to typenames, in their template arguments.

Consider this client-side class:

```
template <typename PointT,
          unsigned short DimensionsCount,
          bool DoBoundsCheck >
struct coord { ... };
```

Note that the last two types in the parameter list are concrete types.

Writing an `s11n_traits` specialization for that is straightforward but tedious. Here's a simpler, more maintainable way to do it:

```
#define S11N_TEMPLATE_TYPE coord
#define S11N_TEMPLATE_TYPE_NAME "coord"
#define S11N_TEMPLATE_TYPE_PROXY s11n::streamable_type_serialization_proxy
#define S11N_TEMPLATE_TYPENAME_T2 unsigned short
#define S11N_TEMPLATE_TYPENAME_T3 bool
#include <s11n.net/s11n/proxy/reg_s11n_traits_template3.hpp>
```

That handles any type of `coord<>`, assuming that `coord<>` is `iostreamable` with the given templated types.

If `coord` is not `iostreamable`, but provides adequate accessor/mutator functions for all of its data, then we can replace the proxy we used above with one specific for `coord` types. For example:

```
struct coord_s11n {
    template <typename NodeT, typename CoordT>
    bool operator()( NodeT & dest, CoordT const & src ) {
        ... copy src data to dest ...
        return true;
    }
};
```

```

    }
    template <typename NodeT, typename CoordT>
    bool operator()( NodeT const & src, CoordT & dest ) {
        ... copy src data to dest ...
        return true;
    }
};

```

Like conventional supermacros, you don't need to undefine the macros you "pass" to `reg_s11n_traits_template3.hpp`. See section [12.1](#) for more general information on supermacros.

The library comes with the following headers to handle such a registration:

```

<s11n.net/s11n/proxy/reg_s11n_traits_template1.hpp>
<s11n.net/s11n/proxy/reg_s11n_traits_template2.hpp>
<s11n.net/s11n/proxy/reg_s11n_traits_template3.hpp>
<s11n.net/s11n/proxy/reg_s11n_traits_template4.hpp>

```

Their usage is identical. If you need a variant for more template arguments, simply use number 4 as a starting point and modify it to suit your needs.

Note that the `S11N_TEMPLATE_TYPENAME_Tn` macros only need to be set if a template argument is a concrete type. For most purposes, the `Tn` macros do not need to be set.

12.8 Where to invoke registration (IMPORTANT)

It is important to understand exactly where the Serializable registration macros need to be, so that you can place them in your code at a point where `s11n` can find them when needed. In general this is very straightforward, but it is easy to miss it.

At any point where a de/serialize operation is requested for type `T` via the `s11n` core framework (including `s11n-lite`), the following conditions must be met:

- The Serializable registration implementation code for `T` must be available to `s11n`. In practice, this means that the registration code must be available to the client code requesting the operation at the time it is compiled.
- `T` must be a complete type, not, e.g. defined only via a forward declaration. (*T's implementation* need not be available, only its *interface* declaration.)

Because of `s11n`'s templated nature, these rules apply *at compile time*. This essentially means that the registration should generally be done in one of the following places:

- `T`'s header file. Most straightforward, but also the sloppiest, as it ties type `T` very closely to `libs11n`. This may also increase compile times noticeably, as it requires including other `s11n`-related headers which might otherwise have no reason to be in your header files.
- The implementation file(s) calling the serialization operation. (Be careful to avoid undue duplication of macro calls, for maintenance reasons and to avoid ODR violations.)
- When Serializables are compiled to standalone DLLs, which is necessarily a polymorphic operation under `s11n`'s classloading model, the class' source file is a good place to put it, as it will only be compiled (and needed) in that one place.
- A separate header created exclusively for this purpose, which is included by any code which initiates de/serialize operations on `T` objects. For example, we might have `T.hpp` and `T_s11n.hpp`, with the latter handling `s11n` registration. This is probably the cleanest solution for non-trivial projects, and is generally the approach taken by `s11n`'s author. It also allows us to ship `T_s11n.hpp` as an "add on" to other code, available for optional inclusion by other `s11n`-using clients.
- In the simplest client-side case, a `main.cpp` with all implementation code in that file, simply call the macros right after each class' declaration. If you later refactor classes out of the main file, move their registration code to one of the places mentioned above.

12.8.1 Hand-implementing the macro code (IMPORTANT)

Whenever these docs refer to calling a certain macro, what they *really* imply is: include code which is functionally similar to that generated by the published macro. This code can be hand-written (and may need to be for some unusual cases), generated via a script, or whatever. In any case, it must be available when `s11n` needs it, as described above.

13 Proxies, functors and algorithms

"Politics is for the moment, an equation is for eternity."

Albert Einstein

s11n's proxying feature is probably its most powerful capability. s11n's core uses it to proxy the core de/serialize calls between, e.g. `FooClass::save_state()` and `OtherClass::operator()`.

Note that any non-serializable type which s11n proxies is actually a `Serializable` for all purposes in s11n. Thus, when these docs refer to a `Serializable` type, they also imply any proxied types. The *proxies*, on the other hand, are not technically `Serializables`.

How to register a type as a proxy is explained in section 12.6.

Most of the classes/functions listed in the sections below live in one of the following header files:

```
<s11n.net/s11n/algo.hpp>
<s11n.net/s11n/proxy/listish.hpp>
<s11n.net/s11n/proxy/mapish.hpp>
```

The whole library, with the unfortunate exception of the `Serializer` lexers, is based upon the STL, so experienced STL coders should have no trouble coming up with their own utility functors and algorithms for use with s11n. (Please submit them back to this project for inclusion in the mainstream releases!)

It must be stressed there is *nothing at all* special or "sacred" about the algorithms and proxies supplied with this library. That is, clients are free to implement their own proxies and algorithms, completely ignoring any provided by this library. If you want, for example, a particular `list<T>` specialization to have a special proxy, that can be done.

13.1 Commonly-used Proxies

This section briefly lists some of the available proxies which are often useful for common tasks.

To install any of these proxies for one your types, simply do this:

```
#define S11N_TYPE MyType
#define S11N_TYPE_NAME 'MyType'
#define S11N_SERIALIZE_FUNCTOR serialize_proxy
// #define S11N_DESERIALIZE_FUNCTOR deserialize_proxy
// ^^^^ not required unless noted by the proxy's docs.
#include <s11n.net/s11n/reg_s11n_traits.hpp>
```

When writing proxies, remember that it is perfectly okay for proxies to hand work off to each other - they may be chained to use several "small" serializers to deal with more complex types. As an example, the `pair_serializable_proxy` can be used to serialize each element of any map. If you write any generic proxies or algorithms which are compatible with this framework, *please submit them to us!*

13.1.1 I/OStreamable types: `s11n::streamable_type_serialization_proxy`

This proxy can handle any `Streamable` type, treating it as a single `Serializable` object. Thus an `int` or `float` will be stored in its own node. While this is definitely not space-efficient for small types, it allows some very flexible algorithms to be written based off of this functor, because PODs registered with this proxy can be treated as full-fledged `Serializables`.

Proxies for the most common PODs come with the library. To register such a proxy, simply do:

```
#include <s11n.net/s11n/proxy/pod/TYPENAME.hpp>
```

If a POD type you are using does not have a proxy header, look at the existing proxies to see how to do this.

13.1.2 Arbitrary list/vector types: `s11n::list::list_serializable_proxy`

This flexible proxy can handle any type of list/vector *containing* `Serializables`. It handles, e.g. `list<int>` and `vector<string*>`, or `list<pair<string, double*>>`, provided the internally-contained parts (like the `pair`) are `Serializable`. Remember, the basic PODs are inherently handled, so there is need to register the contained-in-list type for those or `std::string`.

Registration code for the standard list types can be included like so:

```
#include <s11n.net/s11n/proxy/std/list.hpp>
#include <s11n.net/s11n/proxy/std/vector.hpp>
#include <s11n.net/s11n/proxy/std/set.hpp>
#include <s11n.net/s11n/proxy/std/multiset.hpp>
#include <s11n.net/s11n/proxy/std/deque.hpp>
```

Trivia:

The source code for this type shows an interesting example of how pointer and non-pointer types can be treated identically in template code, including allocation and deallocation of objects in a way which is agnostic of this detail. This makes some formerly difficult cases very straightforward to implement in one function.

13.1.3 Streamable maps: `s11n::map::streamable_map_serializable_proxy`

This proxy can serialize any `std::map`-compliant type which contains Streamable types. This include `std::multimap`.

13.1.4 Arbitrary maps: `s11n::map_serializable_proxy`

Like `list_serializable_proxy`, this type can handle maps containing any pointer or reference type which is itself a `Serializable`.

Registration code for the standard map types can be included like so:

```
#include <s11n.net/s11n/proxy/std/map.hpp>
#include <s11n.net/s11n/proxy/std/multimap.hpp>
```

There is one minor caveat to keep in mind regarding the map proxies: during cleanup after a failed deserialization (section 6.2.1), the cleanup routines cannot explicitly clean up the keys of the maps because they are `const`. In the vast majority of the cases, this is no issue at all. It is only a problem when the keys are pointers. In this case, deserialization will create the objects, but the failed-deser cleanup process cannot deallocate them. If you have maps containing keys of a pointer type, you should be certain to catch any deserialization failures involving the map and deallocate the pointers.

13.1.5 Arbitrary pairs: `s11n::map::pair_serializable_proxy`

Like `list_serializable_proxy`, this type can handle pairs containing any pointer or reference type which is itself a `Serializable`.

This proxy can be installed for `std::pair` types with:

```
#include <s11n.net/s11n/proxy/std/pair.hpp>
```

13.2 Commonly-used algorithms, functors and helpers

The list below summarizes some algorithms which often come in handy in client code or when developing `s11n` proxies and algorithms. Please see their API docs for their full details. Please do *not* use one of these without understanding its conventions and restrictions.

More functors and algos are being developed all the time, as-needed, so see the API docs for new ones which might not be in this list.

function() or functor	Short description
<code>s11n::[list,map]::free_[list,map]_entries()</code>	Deallocates list/map entries. <i>Not</i> for nested containers.
<code>s11n::create_child()</code>	Creates a named data node and inserts it into a given parent.
<code>s11n::find_child_by_name()</code>	Finds a sub-node of a node

	using its name as a search criteria.
<code>s11n::cleanup_serializable()</code> and <code>cleanup_ptr()</code>	Use to generically deallocate Serializable objects.
<code>s11n::map::de/serialize_streamable_map()</code>	Do just that. Supports any map containing only iostreamable types.
<code>s11n::map::de/serialize_[map/list/pair]()</code>	De/serialize maps/pairs of Serializables.
<code>s11n::list::de/serialize_streamable_list()</code>	Ditto, for list/vector types.
<code>s11n::object_reference_wrapper</code>	Refer to an object as if it is a reference, regardless of its pointeriness.
<code>s11n::abstract_creator</code>	Consolidates stack/heap allocation into one API.

As of version 1.1.3, each of the list/map/pair algorithms, plus many of the main algorithms, have an equivalent functor with the same name, plus a suffix of `_f` (as in "functor"). e.g., `serialize_map()` == `serialize_map_f`.

13.3 When proxies aren't desired

Often times, installing a proxy for a type which will be s11n'd at only one code-point is simply overkill. There are also cases where proxies *cannot* be used for a given type T because a different proxy has already installed for T: installing two proxies for one type results in an ODR violation.

In many cases we don't need to install proxies in order to be able to use them. When we, as the designers of serialization algorithms, know that our data can be handled without installing a proxy, we can sometimes use available algorithms directly to achieve the same effects:

```
#include <s11n.net/s11n/proxy/listish.hpp> // list-type algos
typedef std::list<std::string> SList;
...
SList mylist;
...
s11n::list::serialize_streamable_list( destnode, mylist );
// ^^^^ no list proxy needed
```

That particular algorithm only supports lists containing iostreamable types, which do not need a proxy. On the other hand, if we do the following, we would be using a proxy for both our list and string types:

```
#include <s11n.net/s11n/proxy/std/list.hpp>
#include <s11n.net/s11n/proxy/pod/string.hpp>
s11n::serialize( destnode, mylist );
// ^^^^ list and string both need a proxy here!
```

Many of the generic proxies provided with the library need to serialize contained members, e.g. all of the "non-streamable" container-related algos, and use `s11n::[de]serialize()` to do so. This means that they will indirectly require some form of proxy to be installed for their contained types, or will require that the type to be serialized *be-a* serializable.

13.4 Functor tags

As of version 1.1.3, the library declares a number of empty structs as tags for proxies. This allows the following:

- Clearer understanding of the API, as we can now physically stamp all functors with a category label (or labels). This is one approach to saying "functor X models concept Y," and tagging allows us to turn a concept into more than a note in the documentation.
- A nice side effect of tagging is that when generating API docs/class hierarchy views (e.g., with

Doxygen), all tagged structs get grouped by inheritance (tag type). This makes the library easier to get a mental grip on.

- May allow us to use operator and template overloading to assist in composition of functors. e.g., we can compose two nullary de/ser functors into a single nullary functor returning (f1 && f2). With operator overloads and template metaprogramming, using the tag types to guide the way, we could potentially write (f0 = f1 && f2 && f3) to generate a functor which lazily calls three de/ser operations, applying normal logical-and behaviour if either f1 or f2 fail. This will be easier to implement once the C++ TR1/std::tr1 is out in the wild (it contains much better general functor support than the STL).

See the file `tags.hpp` for the full list of tags and the conventions they imply.

Because the de/serialization API has a narrow set of core functions, and a consistent API amongst them, it is hoped that we can create some s11n-specific compositions without having to include a full-fledged composition framework like one provided by Boost.

14 Data Formats (Serializers)

"...control is a degree of inhibition, and a system which is perfectly inhibited is completely frozen."

Alan W. Watts, The Book

That quote might seem a bit out of place, but it is justified: the format of a data file is one way of imposing control over the data. Indeed, all stored data is stored in *some* format or other. In projects which support a single data format (or small number of them), it is not uncommon for the format itself to become a limiting factor in the project's development at some point. That's just plain wrong vis-a-vis modern development techniques, and we will have none of it. One of s11n's goals is to free clients from the restriction of a single format, or even a pair of formats, so that the selection of a data format becomes a background detail, as opposed to a major design decision. In addition to shipping with support for several data formats, users are free to add their own formats on top of the core library.

Ignorance of data formats is all fine and good, but having a serialization library which doesn't ship with support for any formats at all is nearly useless. This section covers the `s11n::io` layer, which is the "default" i/o implementation for the library.

The `s11n::io` namespace provides an interface, generically known as the Serializer interface, which defines how client code initializes a load or save request but specifies nothing about data *formats*. Indeed, the i/o layer of s11n is implemented on top of the core serialization API, which was written before the i/o layer was, and the core is 100% independent of the `s11n::io` layer.

14.1 General conventions

However data-format agnostic s11n may be, all supported data formats have a similar logical construction. The basic conventions for data formats compatible with the s11n model are:

- Each data file contains exactly one root node, per long-standing DOM conventions.
- Nodes may represent any Serializable type, with all that that implies, or "raw" S11n Nodes nodes (storing information not strictly associated with a specific C++ class).
- Nodes may contain an arbitrary number of child nodes.
- Nodes must have a name meeting the criteria specified in section 5.3. The name need not be unique within that branch of the tree.
- Nodes must have an "implementation class name" property set to the class name of the type for which the node contains data. This is used by the classloader when deserializing the node. It is acceptable to use "dummy names" here, provided someone knows how to handle the data without knowing its class name (e.g. the functions described in in section 10.4 work this way). In this library we use `s11n::node_traits<NodeType>::class_name(node)` to set the class name of a node.
- Nodes may contain an arbitrary number of key/value pairs, called properties:
 - Property keys must be unique within any given node, and "should" contain only alpha-numeric characters or underscores, for compatibility with the widest variety of i/o formats. See section 5.3 for the general guidelines.
 - Property values may be of any Streamable Type (*not* pointers) which supports de/serialization via the standard C++ `istream>>` and `ostream<<` operators.

All that is basically saying is, the framework expects that data can be structured similarly to an XML DOM. Practice implies that the vast majority of data can be easily structured this way, or can at least be structured in

a way which is convertible to a DOM. Whether it is an efficient model for a given data set is another question entirely, of course.

14.1.1 File extensions

File extensions are irrelevant for the library - client files may be named however clients wish. Clients are of course free to implement their own extension-to-format or extension-to-class conventions. (i tend to use the file extension `.s11n`, because that's really what the files are holding - data for the s11n framework.)

14.1.2 Indentation

Most Serializers indent their output to make it more readable for humans. Where appropriate they use hard tabs instead of spaces, to help reduce file sizes. There are plans for offering a toggle for indention, but where exactly this toggle should live is still under consideration. On large data sets indentation can make a significant difference in file size - to the order of 10% of a file's size for data sets containing lots of small data (e.g. integers).

14.1.3 Entity translation

Many (most) i/o formats supported by s11n require some form of string translations in order to store data which might otherwise be confused as part of their grammars. These translations happen transparently to users, but it is useful to know about them because:

- You may want to hand-edit your data, in which case you need to ensure that you properly "escape" (translate) your data.
- You might want to save data which has subtle incompatibilities with certain formats.

The translations done by each Serializer are defined in the API documentation for the Serializer class.

As an example of the second point, let's consider that we are saving the raw string `"<<>>"`. Most of you will recognize those characters from XML, HTML, or the like. That string will almost certainly cause problem in the XML-related Serializers, not at serialization-time, but at deserialization-time. The reason is because it may go through the following transformations (depending on the context and the parser, but this is a worst-case):

```
Serialize == "&lt;&lt;&gt;&gt;"
```

```
Deserialize == "<<>>"
```

That deserialized result is certainly not what we saved!

This particular problem is only likely to arise when storing text for use in higher-level parsers, e.g. HTML, and will not happen when storing numbers, simple strings, and the like. The generic translation code has proven to work rather well over the past 1.5+ years, but may get confused in some unusual cases. If you find specific errors, please report them to us (and send us the data file, if possible).

So, though the library is format-agnostic, its users probably should not be. Of the current Serializers, only `compact` does *no* translations, which makes it suitable for use as a data format in cases where the user is concerned about any sort of translation-related mangling. (However, that format is also the least human-readable and not easily hand-editable.)

14.1.4 Magic Cookies

This information is mainly of interest to parser writers and people who want to hand-edit serialized data or generate it from non-libs11n sources, like Perl scripts.

Each Serializer has an associated "magic cookie" string, represented as the first line of an s11n data file. In the examples shown in the following sections the magic cookie is shown as the first line of the sample data. This string should be in the first line of a serialized file so the data readers can tell, without trying to parse the whole thing, which parser is associated with a file. The input parsers themselves do not use the cookie, but it is required by code which maps cookies to parsers. This is a crucial detail for loading data without having to know the data format in advance. (Tip: it uses

```
s11n::cl::classload<SomeSerializerInterfaceType>(first_line_of_input_stream)).
```

Note that the i/o classes include this cookie in their output, so clients need not normally even know the cookie exists - they are mentioned here mainly for the benefit of those writing parsers, so they know how the framework knows to select their format's parser, or for those who wish to hand-edit s11n data files.

Be aware that s11n consumes the magic cookie while analyzing an input stream, so the input parsers do not

get their own cookie. This has one minor down-side - the same Serializers cannot easily support multiple cookies (e.g. different versions). However, it makes the streaming simpler internally by avoiding the need to buffer the whole input stream before passing it on.

See `s11n/io/serializers.hpp` for the API for adding new Serializers to the framework.

Versions 0.9.7 and higher support a special cookie which can be used to load arbitrary Serializers without having to pre-register them. If the first line of a file looks like this:

```
#s11n::io::serializer ClassName
```

then `ClassName` is classloaded as a `Serializer` (a subtype of `s11n::io::data_node_serializer<>`) and, if successful, that object is used to parse the remainder of the stream. Versions 1.1.0+ supports an additional form, functionally identical to the above:

```
#!/s11n/io/serializer ClassName
```

14.2 Overview of available Serializers

This section briefly describes the various data formats which the included Serializers support. The exact data format you use for a given project will depend on many factors. Clients are free to write their own i/o support, and need not depend on the interfaces provided with `s11n`.

Basic compatibility tests are run on the various de/serializers, and currently they all seem to be equally compatible for "normal" serialization needs (that is, the things i've used it for so far). Any known or potential problems with specific parsers are listed in their descriptions. No significant cross-format incompatibilities are known to exist, with the exception that the `expat_serializer` is XML-standards compliant, and is very unforgiving about things like numeric node names.

In some versions of `s11n` the available Serializers are shipped as DLLs, not linked in directly with the library. In these environment, `s11n` tries to auto-load the "known" Serializers (those described below) at startup, but clients will have to load their own DLLs if they have custom serializers. See the `s11n::plugin` API and the existing Serializers for how this is done.

14.2.1 compact (aka, 51191011)

Serializer class: `s11n::io::compact_serializer`

This Serializer read and writes a compact, almost-binary grammar. Despite its name (and the initial expectations), it is not always the most compact of the formats. The internal "dumb numbers" nature of this Serializer, with very little context-dependency to screw things up while parsing, should make it suitable for just about any data.

Known limitations:

- Hand-editing it is very difficult. The data's sizes are encoded in the stream, preceeding the data, and any change in the data requires an update to the size - failing to do so effectively corrupts the data.
- Node/key/class names are limited to 255 characters.
- Property data is "limited" to 4GB per property.

Sample:

```
5119101136  
f108somenode06NoClasse101a0003foo...
```

14.2.2 expatxml

Serializer class: `s11n::io::expat_serializer`

This Serializer, added in version 0.9.2, uses `libexpat`³⁷ and is only enabled if the build process finds `libexpat` on your system. It is grammatically similar to `funxml` (section 14.2.4), but "should" be more robust because it uses a well-established XML parser. Additionally, it handles self-closing nodes, something which `funxml` does not do.

Known limitations/caveats:

- Does only very rudimentary character translation for XML entities - just enough for the input parser to

³⁶ "5119" is as close to "s11n" as i could get with integers. "1011" represents the data format version (there was a predecessor in 0.6.x and earlier).

³⁷ <http://expat.sourceforge.net>

- reliably handle it. This will be fixed when problematic data actually shows up in a use-case.
- Not thread-safe: it is not safe to read from more than one of these objects at a time, e.g. in a client/server environment.
- XML standards compliant, which means it does not tolerate extensions supported by the other s11n XML formats, like numeric node names.

Sample:

```
<!DOCTYPE s11n::io::expat_serializer>
  <nodename class='SomeClass'>
    <property_name>property value</property_name>
    <prop2>value</prop2>
    <empty_property/>
    <empty_class class='Foo' />
  </nodename>
```

14.2.3 funtxt (aka, SerialTree 1)

Serializer class: `s11n::io::funtxt_serializer`

This is a simple-grammared, text-based format which looks similar to conventional config files, but with some important differences to support deserialization of more complex data types.

This format was adopted from libFunUtil, as it has been used in the QUB project since mid-2000, and should be read-compatible with that project's parser. It has a very long track record in the QUB project and can be recommended for a wide variety of common uses. It also has the benefit of being one of the most human-readable/editable of the formats.

Known caveats/limitations:

- Known to have problems reading some unusual string constructs, such as properties which start with a quote but do not end with one.

Sample:

```
#SerialTree 1
nodename class=SomeClass {
  property_name property value
  prop2 property values can \
    span lines.
  # comment line.
  child_node class=AnotherClass {
    ... properties ...
  }
}
```

Unlike most of the parsers, this one is rather picky about some of the control tokens³⁸.

- Closing braces must be on a line by themselves.
- Each property must be on its own line, but may span lines if each newline is backslash-escaped. Such newlines are retained when the data is read in.

This parser accepts some constructs which the original (libFunUtil) parser does not, such as C-style comment blocks, but those extensions are not documented because i prefer to maintain data compatibility with libFunUtil, and they play no role in the automated usage of the parser (they are useful for people who hand-edit the files, though).

14.2.4 funxml (aka, SerialTree XML)

Serializer class: `s11n::io::funxml_serializer`

The so-called funxml format is, like funtxt, adopted from libFunUtil and has a long track-record. This file format is highly recommended, primarily because of its long history in the QUB project, and it easily handles a wide variety of complex data.

Known limitations/caveats:

- Does only very rudimentary character translation for XML entities - just enough for the input parser to

³⁸ Hey, it was my first lexer - gimme a break ;) Also, i wanted it to be compatible with libFunUtil's.

- reliably handle it. This will be fixed when problematic data actually shows up in a use-case.
- To help support the various container serialization functions (section 10.4), this parser accepts node names which are numeric. That feature is not compatible with XML standards, and data files which use this feature may not be loadable by most XML tools without some filtering.
- Does not parse self-closing elements, e.g. `<node ... />`.

Sample:

```
<!DOCTYPE SerialTree>
<nodename class='SomeClass'>
  <property_name>property value</property_name>
  <prop2>value</prop2>
  <empty></empty>
</nodename>
```

14.2.5 parens

Serializer class: `s11n::io::parens_serializer`

This serializer uses a compact lisp-like grammar which produces smaller files than the other Serializers in most contexts. It is arguably as easy to hand-edit as `funtxt` (section 14.2.3) and has some extra features specifically to help support hand-editing. It is arguably the best-suited of the available Serializers for simple data, like numbers and simple strings, because of its grammatic compactness and human-readability.

Known limitations:

- No particular problems known.

Sample:

```
(s11n::parens)
nodename=(ClassName
  (property_name value may be a \('non-trivial'\) string.)
  (prop2 prop2)
  subnode=(SomeClass (some_property value))
  (* Comment block.
  subnode=(NodeClass (prop value))
  Comment blocks cannot be used in property values,
  but may be used in class blocks (outside of a property)
  or in the global scope, outside the root node (but after
  the magic cookie).
  *)
)
```

This format generally does not care about extraneous whitespaces. The exception is *property values*, where leading whitespace is removed but internal and trailing whitespace are kept intact.

When hand-editing, be sure that any closing parenthesis [some people call them braces] in property values are backslash-escaped:

```
(prop_name contains a \) but that's okay as long as it's escaped.)
```

Opening parens may optionally be escaped: this is to help out Emacs, which gets out-of-sync in terms of indentation and paren-matching when only the closing parens are escaped. When saving data the Serializer will escape both opening and closing parens.

Historical speculation: that might explain why, in STL documentation, they denote iterator begin/end ranges in the form $[B, E)$, where " $[$ " means inclusive and " $)$ " means exclusive. If the symbols were defined the other way around, such that $(B, E]$ had the same meaning as above, emacs's paren-matching and indentation modes would get out of sync, which would most certainly have frustrated the designers of the STL. :) Even if that is not the case - which it is probably is not - the paren serializer does explicitly have this escaping behaviour to accomodate emacs. Yeah, i know that a real, die-hard, lisp-loving emacs user [with way too much extra energy] would have simply implemented `paren-serializer-mode...` and probably would have implemented the C++-side serializer class on top of it. And it would work, too, because emacs is just cool that way. But i haven't got that much energy, and thus the above-mentioned backslash hack was introduced.

14.2.6 sqlite3

Serializer class: `s11n::io::sqlite3::sqlite3_serializer`

Available as a separate add-on since 1.2.1, this Serializer uses `sqlite3` (<http://sqlite.org>) databases as the underlying data store. It requires `sqlite3`, and won't work with `sqlite 2.x` and earlier.

Once a Serializable is stored this way, it can be queried and updated using SQL via a tool like the `sqlite3` command-line client, or programmatically via one of the `sqlite3` APIs (e.g., C++ or TCL).

Caveat: unlike most Serializers, this type is schizophrenic when it comes to streams, as described below. It is also extremely slow when reading large data sets.

When writing to streams, it writes SQL code, as the database layer cannot work with streams. When writing files, it writes databases. Likewise, when loading from streams it expects SQL. Loading from files can handle either SQL or databases.

This add-on has two different ways to plug in to `s11nlite`:

First, via the normal "magic cookie" method. This is the conventional approach to register Serializers, and needs no special code by clients. The Serializer is registered when you dynamically link against it or compile it directly into your app.

Secondly, we can plug it in using `s11nlite::instance()`, like so:

```
#include <s11n.net/s11n/s11nlite.hpp>
#include <s11n.net/s11n/io/sqlite3/s11nlite_api.hpp>
...
s11n::io::sqlite3::s11nlite_api slite;
s11nlite::instance( &slite );
```

With that in place, calls to `s11nlite::save/load()` will be routed to the given API handler. In this case, the handler uses the DB-based Serializer for all saving, and automatically determines a Serializer for loading (i.e., it works almost just as standard `s11nlite` when *loading*, with the addition of being able to open database files).

We must ensure that the `s11nlite_api` object outlives all calls to `s11nlite::load/save()`, or that we call `s11nlite::instance(0)` to restore the default handler.

A short demonstration of how to use it with `s11nconvert` and `sqlite3`:

```
stephan@owl:~/> s11nconvert -f foo.s11n -dl sqlite3_serializer -s sqlite3 -o
foo.db
stephan@owl:~/> sqlite3 foo.db
SQLite version 3.2.7
Enter ".help" for instructions
sqlite> .tables
p n
sqlite> select count(*) from p;
308
sqlite> select count(*) from n;
475
```

It can be plugged in to `s11nbrowser` via the Open DLL dialog or the `-dl` command line option. `s11nbrowser` versions newer than 25.11.2005 should work with this plugin.

14.2.7 simplexml

Serializer class: `s11n::io::simplexml_serializer`

This simple XML dialect is similar to `funxml`, but stores nodes' properties as XML attributes instead of as elements. This leads to much smaller output but is not suitable for data which are too complex to be used as XML attributes.

This format handles XML CDATA as follows:

- Only CDATA wrapped in `<![CDATA[a block like this]]>` are recognized.
- At input-time all XML CDATA is stuffed into the "CDATA" property of the node.
- At output-time any data in a node's CDATA property is *not* saved as an XML attribute named "CDATA", but is instead stored as an XML CDATA block.

This is a non-standard extension to data node conventions, so clients which rely on this feature will be dependent on this specific Serializer. (Historical note: i wrote this Serializer in October, 2003, and have never once used the CDATA feature outside of test cases.)

Known limitations:

- See the caveats/limitations notes in section 14.2.4. Most of those apply here.
- Not suitable for use with data which cannot be safely stored as XML attributes. That is, it is fine for storing numbers and other simple types, but storing complex strings may result in Grief (in the form of un-readable data).
- The XML attribute name "s11n_class" is reserved for use by the Serializer in storing each node's class name.

Sample:

```
<!DOCTYPE s11n::simplexml>
<nodename s11n_class='SomeClass'
  property_name='property value'
  prop2='&quot;quotes&quot; get translated'
  prop3='value'>
  <![CDATA[ optional CDATA stuff ]]>
  <subnode s11n_class='Whatever' name='sub1' />
  <subnode s11n_class='Whatever' name='sub2' />
</nodename>
```

14.2.8 wesnoth

Serializer class: `s11n::io::wesnoth_serializer`

"wesnoth" is a simple text format based off of the custom data format used in the game *The Battle for Wesnoth* (www.wesnoth.org).

Caveats and known limitations:

- New (added in 0.9.14) and not well-tested.
- Does not yet properly support multi-line strings as property data. (At least, it's not tested.)
- Up until version 1.2.3, it had a horrible bug which caused many property names to get mangled during deserialization.

Sample:

```
#s11n::io::wesnoth_serializer
[s11nlite_config=s11n::data_node]
GenericWorkspace_size=1066x858
s11nbrowser_size=914x560
serializer_class=wesnoth
[/s11nlite_config]
```

14.3 Tricks

14.3.1 Using a specific Serializer

Easy: simply pick the Serializer class you would like and use its `de/serialize()` member functions. Rather than including its headers, you can load it dynamically:

```
s11nlite::serializer_interface * serializer =
    s11nlite::create_serializer( 'parens' );
```

Normally you *must* select a class (i.e., file format) when *saving*, but loading can be done transparently of the format for the vast majority of cases.

14.3.2 Selecting a Serializer class in s11nlite

See `create_serializer(string)`, which takes a classname and can load any registered subclass of `s11nlite::serializer_interface`. Alternately, set the framework's default serializer type by calling `s11nlite::serializer_class(string)`. As of 1.1, this setting is no longer automatically persistent

across all s11n clients: client applications must either set this at some point or rely on the compiled-in default (which will be some built-in Serializer, but *which one* is not specified by s11n-lite's interface).

14.3.3 Multiplexing Serializers

This has never been done, but it seems reasonable:

If you'd like to save to multiple output formats at once, or add debugging, accounting, or logging info to a Serializer, this is straightforward to do. As of 1.1.2, you can achieve this by subclassing `s11n-lite::client_api<NodeType>` and calling `s11n-lite::instance()`. The "hard core" way to do it would be create a Serializer. By subclassing an existing Serializer it is straightforward to add your own code and pass the call on.

Saving to multiple formats is only straightforward when the serializer is passed a filename (as opposed to a stream). In this case it can simply invoke the Serializers it wishes, in order, sending the output to a different file. Packaging the output in the same output *stream* is only useful if this theoretical Serializer can also separate them later. i can personally see little benefit in doing so, however (maybe a more creative soul can find a clever use for it, though... e.g. protocol-within-protocol wrapping for an RPC channel).

14.4 Internals: flex's role in s11n

This section is intended only for those interested in the implementations of most of the current Serializers. It will be of no interest to anyone else.

The following Serializers have input parsers written using the ubiquitous GNU Flex tool. While it is a powerful tool, its use in modern C++ projects introduces a couple challenges:

- It generates C code. It can be told to output C++ code, but this has problems of its own, not the least of which is the shortage of documentation and its "experimental" status since the late 90's.
- Flex-generated C++ code will not compile as-is under modern compilers because of stricter standards support in today's tools. More recent versions of flex, posted on SourceForge, generate non-compilable code as well, but in other ways. (There's a good reason most Linux distros are still shipping 2.5.4.)
- It is difficult to introduce more than one flex-based parser into a project. The lexer subclassing technique is macro-based, and this ends up causing no end of grief when mixing parsers in a projects. This is particularly troublesome in combination with templates (which are normally inlined in headers).
- The lexer code has to be generated on a system with flex. This rules out most Win32 systems immediately. Even on Unix systems, the generated code won't compile as-is on newer compilers and has to be patched up with perl or sed before compiling it. While this type of manipulation is easy enough to integrate into Unix-based Makefiles, it is not at all trivial for most Win32 environments.

i am not *proud* of the fact that the parsers are built on top of flex. When starting out writing parsers, it was the only tool i knew about, so i used it. And flex is still, after all these years, the only tool of its kind which is well-distributed amongst Unix systems.

The main reasons that most of the Serializers are still implemented in flex, as opposed to re-implementing them in something more modern, are, in order of priority:

1. i am so *damned sick* of writing parsers. i can't look at another one for a while. If *you* want to do it, i would be grateful.
2. There is no other "universally available" parsing kit for C++ out there. There are lots of projects who aspire to do this, but many are commercial, and various ambitious Open Source projects of this type have petered out without producing a usable product.
3. The s11n source tree has a good deal of underlying support code (both C++ and Makefile rules) to integrate flex-based parsers into the library, such that they can be built as "built-ins" or dynamically loaded without the library caring which they are. That code's been around a long time and works quite well, so i'm in no hurry to replace it. Using that backbone, writing a new flex-based Serializer is normally only a few hours of work.

Long-term, i would eventually like to reimplement the parsers in, e.g., Spirit (<http://spirit.sourceforge.net>), but see point #1 in the above list. Initial experimentation with Spirit suggests that it requires that it buffer all input before tokenization starts. Experience has shown that this is not an acceptable option for this library, as it can drastically affect runtime speed of large data sets, and inherently increases deserialization memory requirements by roughly a factor of one. See section 25.4 for more information on the implications of such a copy.

15 `class_name()` and friends

"A rose by any other name would smell as sweet."

Shakespeare

"But a class not derived from T is-not-a T."

Anonymous Software Developer

Once upon a time - the first few months of s11n's development - s11n developed a rather interesting trick for reliably getting a type's *name* at runtime. Despite how straightforward this must sound, i promise: **it is not**. C++ offers no 100% reliable, in-language, well-understood way of getting something as seemingly trivial as a type's *friggling name*. While s11n's trick (shown soon) works, it has some limitations in terms of cases which it simply cannot catch - the end effect of which being that objects of BType end up getting the class name of their base-most type (e.g. "AType"). Let's not even think about using typeid for class names: `typeid::name()` **officially provides undefined behaviour**, which means we won't even consider it.

Historical note: Very early versions of s11n used a typeid-to-typename mapping, which worked quite well (and did not require consistent typeids across app sessions), but it turns out that `typeid(T).name()` can return different values for T when T is used different code contexts, e.g. in a DLL vs linked in to the main app. Thus that approach was, sadly, abandoned. i even used an external database at one point: dump the symbols from your object files, using nm, into a file and we read those at runtime to get the class names. While remarkably effective, that turned out to be about as fun to maintain as poison ivy is to play in.

To be honest, the details of class names vis-a-vis s11n, in particular vis-a-vis client-side code, are an amazingly long story. We're going to skip over significant amounts of background detail, theory, design philosophy, etc., and cut to the "hows" and the more significant "whys".

15.1 `node_traits<T>::class_name()`

Note: in older s11n code we had an `impl_class()` function. That was identical to `class_name()`, but is long-since deprecated. Some documentation may still refer to `impl_class()` in some cases, but these can be safely understood to mean `class_name()`.

For s11n, a node's metatype class name is significant at the following points:

1. When serializing an object, the node it is stored in should have its `class_name()` set to the object's class name. This is trivial to achieve at the framework level for the majority of (all?) *monomorphic* types, but impossible to achieve *polymorphically* without some small amount of client-side work. In s11n this "small amount" of work comes in the form of setting a node's `class_name()` to the string form of the Serializable's class' name. This is done in an object's serialize operator (not deserialize). If a type inherits Serializable behaviours it must set the `class_name()` **after** calling the inherited behaviour, to avoid that the parent type overwrite the `class_name()` of the subtype. Note that Serializable Proxies need to set the name of the *Serializable type*, **not** to the name of the *proxy type*. Why? Read the next section and then it should be clear.
2. When deserializing a node to a given InterfaceType, as in this code:

```
InterfaceType * b = s11n::deserialize<InterfaceType>( somenode );
```

s11n asks the InterfaceType's classloader for an object of the type mapped to the name stored in `node_traits<NodeType>::class_name(somenode)`. The classloader, ideally, has a subtype of InterfaceType registered with that name (or it is InterfaceType's name, or maybe it can find the type via a DLL lookup). If so, the classloader will return a new instance of that type and s11n will hand off the data node to it using the internal API marshaling interfaces. If no class of the given name can be found by InterfaceType's classloader (other classloaders are not considered), deserialization necessarily fails, as there is no object to deserialize the data into.
When a data node is "directly" handed to a Serializable (e.g. `s11n::deserialize(srcnode, targetserializable)`) then the class name is *irrelevant*, as s11n must assume that the given node and Serializable "belong together", semantically speaking. This property can be used to store arbitrary data in nodes and have a complementary deserialize algorithm or functor which understand the "data layout" within the node. e.g. the various `serialize_streamable_xxx()`

variants use this: each pair of de/serialize functors supports one end of the data's "dialect", would be one way to put it. This can be used to de/serialize some objects which are themselves not registered as Serializables, by simply "walking" them in our algorithm. In fact, in this case the only reason such types cannot be called true Serializables is because s11n's API does not have (is not given) a registered proxy through which to redirect them.

In theory these points are all pretty straightforward, and all should make pretty clear sense. After all, to load a specific type it must have a lookup key of some type, and a classname makes a pretty darned convenient key type for a classloader. The classloader's core actually supports any key type, but s11n is restricted to strings, mainly for the point just mentioned, but also because non-strings aren't meaningful in the context of doing DLL searches for new Serializable types. Consider: what should an int key type be useful for in that context - interpreting it as an inode number? Thus, s11n internally uses only string-keyed classloaders. This is not to say that the string must be the same as a class' name: you may of course use numeric strings.

Hopefully the significance of a node's class name is now fully understood. If not, please suggest how we can improve the above text to make it as straightforward as possible to understand!

Side-notes:

- i do honestly believe it to be impossible in C++, using *only* in-language techniques, to 100% reliably get the class name for polymorphic types, not considering options like external (file-based) lookup tables. ***i would be extremely happy to be proven wrong!*** Please contact the development mailing list if you know a magic trick for this!
- s11n actually did use external lookup tables for class names once, created by using the nm tool to extract all type names from an application/DLL *after* linking it. The immediate advantage is that it works fairly well, as it has access to all class names used in the binary (app/DLL), but it's cumbersome, build-wise, and *very* memory-hungry, as a huge number of the types in any binary are not at all relevant to the client for purposes of s11n (e.g. `std::__gcc_blahblah_internal<Foo` *, `std::allocator<Foo>`), and we can't reliably programmatically determine what we could cut out).

15.2 s11n_traits<T>::class_name(const T *)

In s11n 1.1.0, `s11n_traits` was expanded to replace the former `class_name<>` type (and the scattered kludges which cropped up around it).

Many of the shipped algorithms use this API to get node's class name, as described in the previous section.

Clients who have types which have a function allowing them to return their real class name can specialize `s11n_traits` for their type to allow s11n to internally get access to the proper class names. An example specialization of this function might look like:

```
static std::string class_name( const T * hint ) {
    return (!hint)
        ? 'T' // return Interface Type's name
        : hint->className(); // assuming T's API has such a feature
}
```

15.3 Class name of "unknown"

Sometimes you may see a class name of "unknown" in your data. This is not necessarily a problem, and can be caused by the following:

```
typedef std::list<std::string> SL;
SL li;
... populate li ...
s11n::list::serialize_streamable_list( destnode, li );
```

Algorithms get their type's name by using `s11n_traits<T>`, and in the above case there isn't necessarily an `s11n_traits<SL>` installed because the list type was never explicitly registered as a Serializable (it doesn't need be to for this case).

This is actually all fine and good, and will not cause any problems in a case like the one above. If you desperately want to set a class name, it is okay to manually do so in a case like this (but not as a general rule: see section 23.5.1).

In fact, for all deserialization which does not involve pointers, the logical classname of a node *is ignored*., as the s11n'd data is fed to pre-existing objects. In the case of pointers, we use the classname to load the object

and then pass that object through the deserialization process just as we do any non-dynamically-allocated object.

16 Exceptions conventions

"I need a woman who can say, 'honey, can you please take a look at this stack trace while I order the pizza?' and really mean it."

Anonymous Software Developer

Please also see the section 19, which is closely related to this material.

As of version 1.1, s11n attempts to define a set of exception-related guarantees, such that we can define the state of, e.g. a container, when the de/serialization of a child node fails.

It is important to always remember that, like most other software, s11n requires that destructors never throw. If a dtor throws then all exception guarantees go out the window. Likewise, if a *default* ctor or a copy/assignment ctor throws, guarantees may go bye-bye.

The base-most exception type for the framework is, naturally enough, `s11n::s11n_exception`, which derives from `std::exception` and follows the same interface. The API does not have any `throw(xxx)` specifiers on most functions. This is to allow the library to propagate user-thrown exceptions without running the risk of `unexpected()` being called (that's C++'s way of crapping out if a function throws an exception which does not match its `throw(xxx)` specification). All functions in the API should accommodate the propagation of exceptions, preferably with well-defined results. The exact guarantees regarding any throw behaviour are necessarily documented on a per-algorithm basis, so see the appropriate API docs. Almost all recursive routines go through the core de/serialize and may throw, but the exact definition of what happens in the face of exceptions must be defined by each algorithm.

Note that no amount of conventions will 100% transparently protect clients from problems such as memory leaks. As of version 1.1.3, the library is believed to be able to protect from all leaks it possibly can. It has no known leaks in valid use cases, and allows clients to extend the cleanup support such that their types can be guaranteed not to leak if a deserialization fails, whether it fails due to an exception or not.

16.1 The library throws when...

The core library itself never throws. It will pass on exceptions, but it does not throw any simply because all the real work is delegated.

The various layers built around the core may or may not throw. The guidelines are:

- The support algorithms, like the containers proxies, may throw whenever they like. Precondition violations are prime candidates for throwing.
- The plugins layer does not throw, but currently only due to dependencies reasons, and this may change at some point.
- The i/o layer may throw exceptions whenever it likes.
- `s11n_lite` is primarily a wrapper, and may propagate exceptions passed on through the core, plugins, i/o, or `client_api` layers.
- The post-failure cleanup support explicitly catches and discards all exceptions, to ensure no-throw-on-destruction semantics.

Plugin operations are called during the deserialization process to find unknown types. In theory they may throw, but they currently do not. This no-throw policy is under consideration, and likely to change at some point.

16.2 Throwing from client-side de/ser operations

Let's consider the following deserialization operator for class ST:

```
bool operator()( const s11n_lite::node_type & src ) {
    typedef s11n_lite::node_traits TR;
    if( ! TR::is_set( src, "some_key" ) ) {
        // this is an error in our case
    }
    ...
}
```

```
}
```

The client has at least three options for how to handle the error:

1. Recover from the error, if possible/desirable. For example, use a default value for the missing data.
2. Return false.
3. Throw an exception.

Options 1 and 2 have been around since the beginning of `libs11n`, but option 3 was introduced in 1.1.0. When a client-side de/serialization algorithm throws, how the internals of the library react to it depends on a number of factors. As of 1.1.3, the major algorithms were reimplemented to deallocate resources properly on exceptions, using `s11n_traits::cleanup_functor` (section 6.2.1). Each algorithm documents its exact behaviour, but the general overall guaranty is that no memory will go leaked if a deserialization fails. In older library versions, this was only true as long as the types which failed to deserialize managed their own memory (i.e., *not* standard containers of pointers, though these are now safely handled).

As a rule, if deserialization of an object fails (returns false or throws), the object is either unmodified (only possible in a few cases) or in an undefined state (the majority of cases). A general prerequisites for when we can apply the non-modified guaranty to a `Serializable` type are:

- A custom algo must be used for the type, or an existing algo with this guaranty must be used. e.g., the default proxies for `std::set` and `std::list` use the same deserialization algo, which happens to provide this guaranty, thus both of those containers provide it when the default algos/proxies are used.
- The type must support an efficient `swap()` feature, or something semantically similar. This is because one of the simplest, most effective, and most efficient ways to implement this guaranty is by using `swap()` after deserialization into an intermediary object succeeds.

In fact, this library could theoretically offer the unmodified guaranty in even the default-most algorithms, for all types, but this would require that all supported types be copyable via the default C++ copy mechanisms, which might not be realistic. It also would not be as inherently generic and efficient as `swap()`. I have reservations against relying on `std::swap()` as the default behaviour because it does not *guaranty* an efficient swap, it only provides a standardized interface for the swap feature. Falling back to `std::swap()` by default would be misleading at best, and may result in unacceptable behaviours in some cases unless `swap()` is reimplemented/overloaded.

16.3 Errors and `SerT * deserialize<NodeT, SerT>(const NodeT &)`

Consider this perfectly innocent-looking call:

```
T * t = s11nlite::deserialize<T>( mynode );
```

What that does is essentially this:

1. Try to instantiate an object of type `node_traits<>::class_name(mynode)`, whose interface `Type` is `T`. If it fails, we can safely signal an error at that point.
2. Calls `deserialize(mynode, *theNewObject)`. If it succeeds, return `theNewObject`. If it fails...

Now the correctness of its behaviour is `T`-dependent. It was not until going over the exceptions support that the inherent danger of deleting the failed object became apparent. Client-written classes normally manage their contained objects' memory, so these are not a problem, but any standard container containing pointers is a problem. If we delete a container object which itself contains pointers or contains, somewhere nested in its subcomponents, any unmanaged pointers (not owned by their containing object) a deletion of `theNewObject` will cause a memory leak.

The `s11n_traits::cleanup_functor` convention was developed to create a safe way for deserialize algorithms to handle such an error case. If an exception is thrown from `deserialize()`, or deserialization otherwise fails, the internally-allocated object can be safely cleaned up via the cleanup functor. For example, all of the following types will be clean up properly in the face of errors, assuming that an appropriate cleanup functor has been defined for each:

```
list<T>
list<int>
map<int, list<multimap<double, T *>>>
```

(The library's default proxies for these types install working clean-up functors.)

See section 6.2.1 for how this works.

16.4 Exceptions and "external modules"

i recently (July 2005) bought the book *C++ Coding Standards*, by Herb Sutter and Andrei Alexandrescu. Item 62 in the book is entitled "Don't allow exceptions to propagate across module boundaries," and explains that, for example, throwing an exception from a de/serialization algorithm is not *actually* guaranteed to be safe if the exception "crosses module boundaries." That is basically to say, thrown from different libraries linked in the same application. Since s11n is implemented largely in header files, those parts which would throw would actually throw from *your* module, because they are compiled as part of your code. There are a few non-template places which can throw as well. Going the other direction: if your class' de/serialization operator throws, that exception must go back through the s11n core before being passed back to the caller. That would normally be fine, but if the class which threw the exception is from another module, it *might* not be possible for your C++ runtime environment to pass the exception from the algo to s11n's core. These types of problems are related to much lower-level operating system and hardware details than the C++ standard can accommodate, and thus the implementation depends 100% on your compiler, linker, and the benevolence of your chosen god(s).

That said...

In practice, it is possible to throw across module boundaries when the throwing module and the modules the error passes through are compiled "using the same options", though what that really means is rather blurry. If, however, you compile library A on compiler version 1.0 and then another module under compiler version 1.2, the results might not be binary-compatible enough to pass exceptions between the two. Again, vendor-dependent.

Considering that i've been using s11n for almost two years now without an exception causing this level of crash, i personally consider this problem to of little concern. Then again, during most of that time, exceptions were explicitly not handled by the library (well, at least not properly), so they were never intentionally thrown during de/serialization. Since 1.1.x it is legal to throw, so... pay heed to the above advice.

16.5 Specific guarantees

The core algorithms *cannot* provide a specific guaranty on the state of an object on which deserialization fails, but as of 1.1.3 many of the major support algorithms can. By extension, this means that using a Serializable type which is handled by these algorithms implicitly gives these guarantees to the core algorithms.

Below is a list of algorithms which provide the following guarantees on a deserialization failure (including exceptions) into a Serializable object we will call Target:

1. Target is not modified.
2. Dynamically-allocated resources contained in Target are deallocated via the `s11n::cleanup_serializable()` mechanism. (Section 6.2.1.)
3. All exceptions are propagated back to the caller.

e.g., when calling `serialize(srcnode, mylist)`, the Target for the deserialization is `mylist`.

Without a doubt, the second guaranty is the most significant. The first guaranty has been waived since s11n's earliest days, but recent code reviews and refactorings provided satisfactory solutions to the cleanup problem, which inherently makes the first guaranty easier to implement, in particular for types which support an efficient swap operation.

The algorithms which explicitly support this are:

- `s11n::list::deserialize_list(...)`
- `s11n::map::deserialize_pair(...)`
- `s11n::map::deserialize_map(...)`
- `TargetT * s11n[lite]::deserialize(const NodeType & src)`
(In this case, the T object might be modified, but the client will never get the object if deserialization fails, so the effect is the same.)

Other algorithms might support these guarantees as well - see the API docs for the algorithms used by your de/serialization proxies/implementations.

16.6 Making your Serializables exception-safe

As of 1.1.3, the `s11n::cleanup_serializable()` mechanism (section 6.2.1) is defined to "clean up" objects which fail to deserialize. Originally conceived to clean up standard containers of unmanaged pointers, a small API has grown up around that type which simplifies leak-protection in many deserialization cases. Let's consider the following code, assuming that it is some client-side code other than a de/serialize operator:


```
s11n_lite::micro_api<MyType> micro;
MyType * myObj = micro.load( 'myfile.s11n' );
if( ! myObj ) { ... loading failed! ... }
...
```

That's all fine and good, but let's assume that either an exception is thrown somewhere immediately afterwards, or that you are in fact utterly lazy and do not want to have to manually delete `myObj`. Both cases have the same solution, which is to:

1. Make sure we have a valid cleanup functor installed. For types which manage/own their own internal pointers, the default functor will do the job - we only need to specifically define one for "container-like" types which hold unmanaged pointers.
2. Use `s11n::cleanup_ptr<MyType>` in a manner similar to how we would use `std::auto_ptr<MyType>`.

Now we simply modify the above code to look like this:

```
s11n::cleanup_ptr<MyType> myObj( micro.load( 'myfile.s11n' ) );
if( ! myObj.get() ) { ... loading failed! ... }
```

Now, when `myObj` goes out of scope, `s11n::cleanup_serializable<MyType>()` will be called to take care of the cleanup process. In fact, for types which manage their own pointers, an `auto_ptr<>` will have the exact same effect *for most type*, but we show the `cleanup_ptr<>` approach for demonstration purposes. For example, the following case would not behave as desired with an `auto_ptr<>`:

```
typedef std::list<MyType *> MyList;
s11n_lite::micro_api<MyList> micro;
MyList * mylist = micro.load( 'myfile.s11n' );
...
```

If we simply delete `mylist`, or use an `auto_ptr<>` to delete it, the pointers in `mylist` will leak! Depending on the size of the list and the items it contains, the leak might be small or *huge*. In any case, no leak is acceptable behaviour.

We can clean up any `Serializable` object, regardless of pointeriness, nestedness, etc. with:

```
s11n::cleanup_serializable<FoosInterfaceType>( foo );
```

We don't care if `foo` is a pointer or reference here, and we don't care what subtype of `Foo` it is.

When using *pointers* to `Serializables`, it is often more convenient to use `cleanup_ptr<>`, as demonstrated here:

```
cleanup_ptr<MyList> mylist( micro.load( 'myfile.s11n' ) );
```

When `mylist` goes out of scope, or when `mylist.clean()` is called, or `mylist` is otherwise reassigned, the list is walked and `s11n::cleanup_serializable<MyType>()` is called on each entry in the list. The effect is that the list entries will get destroyed. Afterwards, the `MyList` pointer itself (if it is a pointer) is destroyed. If `MyList` contains another container, e.g., `std::vector<MyType*>`, then that container will be walked recursively - the end effect is the same, regardless of the nesting level. The only requirement is that the contained type have a `s11n_traits<>::cleanup_functor` which is designed to work with that type (again, most objects can use the default or one of the already-supplied implementations).

Keep in mind that `cleanup_ptr<>` is *only* for use with cleaning up registered *Serializables*, and is *not* a general utility class! If used on non-`Serializables`, it will use the default cleanup functor, which might or might not have the desired results for any given type. The proxies for the standard containers install a cleanup handler for their container type, so when proxying standard containers, the hard part will be done for you. In some cases it is essential to write a custom cleanup functor, however. See the example in `src/proxy/reg_list_specializations.hpp` for how this is done.

17 SAM: Serialization API Marshaling layer

"Play it again, Sam!"

Common proverb

Achtung: SAM is not Beginner's Stuff. This is, as Harald Schmidt puts it so well in a German

coffee advertisement, *Chefsache* - intended for use by the "higher ups." This is *not* meant to discourage you from reading it, only to warn you that in s11n-lite, and probably even when using the core directly, you will normally never need to know about SAM. There may be some unusual cases where writing a SAM specialization is just what is needed, however.

Achtung #2: There is a fine line, and indeed some overlap, between certain responsibilities of SAM and those of `s11n_traits<>`... but the line isn't well-defined and the small overlap is actually a flexibility benefit (e.g. where is a node's `class_name()` set?). In effect, `s11n_traits<>` provides the *public* interface for API marshaling and SAM provides the *s11n-internal* interface. Traits and SAM also each have some very distinct responsibilities, and consolidating them into one type is not planned.

It's time to confess to having told a *little white lie*. Repeatedly, even willfully, *many* times over in this span of this document.

The Truth is:

s11n's core doesn't actually implement its own "Default Serializable Interface"

WTF? If s11n doesn't do it, who does?

Following computer science's oft-quoted "another layer of indirection" law, s11n puts several layers of indirection between the de/serialization API and... *itself*. To this end, s11n defines a minimal interface which describes only what the s11n core needs in order to effectively do its work - no more, no less. s11n sends all de/serialize requests through this interface, which is generically known as:

SAM: Serialization API Marshaling³⁹ layer

i admit it: i have, so far, *willfully* glossed *right* over SAM. However, i did so *purely* in the interest of keeping everyone's brains from immediately going all wahooney-shaped when they first open up the s11n manual. As *you've* made this far in the manual, we can only assume that wahooney-shaped brains suit you just fine. If that is indeed the case, keep reading to learn the Truth about SAM...

17.1 The SAM layer & interface

i've been telling you this *whole time* that types which support s11n's *Default Serializable Interface* are... well, "by default, they're already Serializables." In a sense, that's correct, but only in the sense that i've been "abstracting away" the very subtle, yet very powerful, features implied by the existence of SAM. Bear with me through these details, and then you'll surely understand why SAM is buried so far down in the manual.

At the heart of s11n, the core knows only about these small details:

- SAM's two API functions and their conventions (which are identical to those of s11n's core de/serialize functions).
- `node_traits` (section 6.1), and only a small portion is used internally.
- `s11n_traits` (section 6.2).

s11n's core doesn't know *anything* about *anyone's* de/serialize interface *except* for that of SAM's. The core, to be honest, is essentially quite dumb - implemented in a relative *handful* of lines of code - looking over the code now i'd guess that, if we don't count the `[de]serialize_subnode()` convenience funtions, it's *less than 30* actual code lines(!!!).

SAM defines the interface between s11n's core and the world of client-side code. The following code reveals the *entire* client-to-core communication interface:

```
template <typename NodeType, SerializableT>
struct s11n_api_marshaler {
    typedef SerializableT serializable_type;
    typedef NodeType node_type;
    static bool serialize( node_type &dest,
                          const serializable_type & src );
    static bool deserialize( const node_type & src,
                             serializable_type & dest );
};
```

(Prior to 1.1.3, the `NodeType` parameter was a template parameter for the *functions*, but not the *class*. This chapter normally refers to the older signature, but this difference is insignificant for most purposes.)

By now that interface should look eerily familiar. Note that static functions were chosen, instead of functor-style

³⁹ Note that both "marshaling" and "marshalling" are correct spellings of this word. s11n uses the single-l variant because `ispell` told me that was correct ;)

`operator()`s, based on the idea that these operations are activated very often, and i felt that avoiding the cost of such a frivolous functor was worth it. Additionally, this interface defines something "solid" for clients, as opposed to s11n's normal convention of using two overloads of `operator()`. There's another, somewhat lamer, reason the `operator()`-style interface can sometimes cause ambiguity errors, so it needs to be avoided here.

SAM specializations may define additional typedefs and such, but the interface shown above represents the core interface: extensions are completely optional, but reduction in the interface is not allowed.

It is important to understand *how* s11n "selects" a SAM specialization: by the *type* argument passed as a `SerializableType` template parameter. Thus, s11n uses a `SAM<myobject's type>` specialization. We've jumped ahead just a tad, and it's now time to back up a step and, with the above in mind, get a better understanding of SAM's place in the s11n model...

17.2 SAM's place in the API calling chain (and other important notes)

After client code initiates a de/serialization operation, the process goes something like this:

1. s11n passes off the call to `s11n_api_marshaler<T>::[de]serialize(node, obj)`.
2. SAM is now in control of the request. The default SAM implementation simply sets the node's class name, using `s11n_traits<T>::class_name()`, and delegates the request to `s11n_traits<T>::[de]serialize_functor`, as appropriate.
3. SAM eventually returns to the core, which then passes the results directly back to the user.

In API terms, SAM is *the internal place* to manipulate the marshaling process, e.g. to implement custom API translation. The *public* interface for doing so is by specializing `s11n_traits` for a given type.

As a special case⁴⁰, `SAM<X*>` is single implementation, not intended to be further specialized - see below!

Note that in this context, "client code" might actually refer to an algorithm or functor shipped with s11n - as far as the core is concerned anything, including common "convenience" operations (e.g. child node creation), which happen before the the core calls SAM, and while waiting on SAM, are "client code."

17.2.1 More about `SAM<X*>`

A *single* specialization of `SAM<X*>` does pointer-to-reference argument translation (since its `SerializableTypes` will be pointer types) and forwards them on to `SAM<X>` (unless they are 0, in which case it simply returns false - effectively a failed de/serialization attempt). Thus pointers and references to Serializables are internally handled the same way (where practical/possible), as far as the core API is concerned, and both `X` and `(X*)` can normally be used interchangeably for `Serializable` types passed to de/serialize operations.

The end effect is that if a client specializes `SAM<Y>`, calls made via `SAM<Y*>` will end up at the expected place - the client-side specialization of `SAM<Y>`, and the pointer will be dereferenced before passing it to `SAM<Y>`.

Some coders show a level of distrust for this "feature", but practice has shown that it is 100% non-intrusive, 100% predictable, and allows some tricks which are otherwise difficult to achieve. In fact, code related to this specialization has not needed any maintenance since its initial introduction, a bit more than a year ago - it is a pure background detail.

Client code SHOULD NOT implement any pointer-type specializations of `s11n_api_marshaler<X*>`⁴¹. Clients MAY implement such specializations, but they're on their own in that case. As it is, if a client implements a `SAM<X*>` specialization the effects may range from no effect to a very difficult-to-track discrepancy when *some* pointer types aren't passed around the same as others. Then again... maybe that's *exactly* the behaviour you need for type (SpecialT*)... so go right on ahead, just be aware of s11n's default handling of `SAM<X*>`, and the implications of implementing a pointer specialization for a SAM. Such tricks are not recommended, and related problems could be extremely difficult to track down later.

17.3 Historical changes

In 1.1.3, the following significant changes were made to `s11n_api_marshaler<>`:

- `DataNodeType` templated type was moved from the functions to the class, to allow for full client-side

⁴⁰ Now that i re-read this, this is one of *extremely* few "special cases" in s11n. i have a special type of *non-love* for "special cases" in general, and avoid them in the interfaces at all costs.

⁴¹ without much consideration, that is. There are conceivable uses for this, but they seem to be well beyond the realm of "common serialization needs", and thus we won't dwell on them here.

- specialization.
- Moved from an anonymous namespace into the `s11n` namespace. The anonymous namespace appears to be unnecessary, and may never have been necessary. (It was there for a reason, but that was soooo long ago...)

18 s11nlite specifics

"People don't do what they believe in. They just do what's most convenient, then they repent."

Bob Dylan

The `s11nlite` API provides a simplified interface into `s11n`. It is intended to simplify the majority of client-side calls into the core library, primarily by abstracting away the `Data Node Type` which is so prevalent in the core API. The "lite" API also wraps up the `s11n::io` API, so it provides a simpler interface into i/o as well. `s11nlite` is intended for "top-level" client use, whereas the core library is more suitable for implementing the internals of specific de/serialization algorithms.

This section covers `s11nlite`-specific behaviours which are not covered by the core library.

While `s11nlite` is a complete client-side interface into `s11n`, `s11nlite` does *very* little work itself: it mainly forwards calls to the core and i/o layers.

18.1 Why use s11nlite?

(Please also see the notes about `s11nlite` in section 2.5.)

By using `s11nlite` as the main client-side interface, client code can be *significantly* simplified over using the core `s11n` and `s11n::io` APIs directly. The main difference is a *lot* less typing of template types. Also, the benefit of fewer direct dependencies on `s11n`-related types should not be underestimated. A concrete example of these simplifications, compare the following two function signatures:

```
s11n::serialize<s11n::s11n_node,MyType>( destnode, srcobj );
s11nlite::serialize<MyType>( destnode, srcobj );
```

The different might appear trivial, but trust me, the first form gets annoying really quickly.

Actually, in the case of monomorph types and the base-most types in a hierarchy of `Serializable`s, C++'s automatic template type deduction can eliminate the need to be explicit about `MyType` when using the first form. The gotcha is in polymorphism: we need to be sure to base the base-most `MyType` in the hierarchy, so we really *should* be explicit when using the first form, or the proper underlying helper types might not be selected (those associated with the base interface in the hierarchy), which ends up leading to confusing compile errors or potentially runtime errors.

Some developers might recommend swapping the order of the template args in

`s11n::somefunc<NodeType,OtherT>()`, as node types are almost always monomorphic and thus their types can be accurately deduced. That would lead to client-side calls like:

```
s11n::serialize<MyType>( destnode, srcobj );
```

Early versions of `s11n` had this convention, with the `NodeType` always as the trailing arg. As it turns out, always having the node object as the first function argument fits in more consistently in the overall API, and i want the template parameters to be in the same order as the function arguments.

`s11nlite` was primarily developed to simplify this type of detail, but also to provide a link to the i/o layer, as the core is blissfully unaware of the pains of i/o.

18.2 client_api<NodeType>

As of `s11n 1.1.0`, `s11nlite` is based upon a class called `client_api<>`. This was done primarily because experience showed that `s11nlite` was not extendable by clients without literally hacking in their desired features. A short background story, to put this into context:

As an experiment, in late 2004 i hacked together a copy of `s11nlite` which used the network layer of the `P::Classes` project (<http://pclclasses.com>). This allowed saving over ftp, for example. The problem was, clients wishing to use it had to know specifically about it (called `ps11n`), and write to its API, which was the exact same as `s11nlite`'s except for the namespace. The end result was two usage-compatible, data-compatible, but completely independent libraries.

Factoring out the main s11n-lite functionality into a subclassable type provides a solution which allows all s11n-lite client code to stay inter-compatible, even when they each use customized back-ends (i.e., their own `client_api<>` subclass, or one provided by a 3rd party library).

Much of the s11n-lite API internally uses an instance of `client_api<>`, which can be fetched or set via the following functions from the s11n-lite namespace

```
client_interface & instance();  
void instance( client_interface * newinstance );
```

(`client_interface` is a typedef for `client_api<s11n-lite::node_type>`.)

See the API docs for the conventions and rules, in particular the ownership rules for the setter.

This feature allows clients to use the s11n-lite API as a front-end for customized extensions to s11n-lite. Without this support, extending s11n-lite while maintaining cross-client API compatibility at the same time is essentially impossible.

The end result is: by extending `client_api<>`, clients can write custom s11n-lite-like APIs, or s11n-lite-compatible extensions, with very little effort. With a bit of additional effort a client can even support *multiple* back-ends at once, though I honestly can't think of a useful case for this.

18.3 File formats

The lite library likes to hide the detail of file formats from you, but does allow you to specify your preferred format:

```
s11n-lite::serializer_class( 'ClassNameOfSerializer' );
```

This preference stays in effect until set again. Unlike version 1.0, in 1.1+ it is not persistent *across application sessions* because it was simply too annoying to have each app overwrite the default of every other app.

We can create a Serializer of a given class with:

```
s11n-lite::serializer_interface * ser = s11n-lite::create_serializer(  
'ClassName' );
```

This will return 0 on error, and does *not* set the library-wide preference.

The classname passed to these functions must be a string associated with a Serializer class, either built-in or dynamically loadable (if plugins support is enabled in your s11n). Most Serializers are registered under *three* names: their formal name, a convenience name, and their "magic cookie". For example, the following calls all have the same effect:

```
s11n-lite::serializer_class( 's11n::io::funtxt_serializer' ); // formal  
name  
s11n-lite::serializer_class( 'funtxt' ); // convenience name  
s11n-lite::serializer_class( '#SerialTree 1' ); // magic cookie
```

It is not recommended to use the cookies directly in client code. The formal names are more preferred, but convenience names are there for a reason - convenience (especially for use when passing the class names as command-line arguments). By convention, the convenience name is always the class name of the Serializer, stripped of namespace and the `_serializer` suffix (if any).

It is up to each Serializer to initially register any names under which it is available. Registering the cookie is required for dynamic file dispatching to work, but the other names are conventionally registered as well (mainly for potential client-side use).

18.4 Simple config files

s11n 1.1.3 adds the `s11n-lite::simple_config` class. It simply acts as a wrapper for a single s11n node, loading it upon construction and saving it upon destruction. Here is how to use it:

```
#include <s11n.net/s11n/simple_config.hpp>  
...  
s11n-lite::simple_config config('MyApp-1.0');  
using std::string;  
typedef s11n-lite::simple_config::node_traits TR;  
string somestring = TR::get( config.node(), string('somekey'), string() );
```

```
s11nlite::serialize_subnode<MyType>( config.node(), mySerializableObject );
```

The ctor will attempt to load the file `$HOME/.MyApp-1.0.s11n`. If `$HOME` cannot be resolved (via a call to `::getenv()`) then the ctor will throw a `std::runtime_error`. If the internal call to `s11nlite::load_node(...)` fails then we cheerfully assume the file didn't exist and create a new one. The file will be saved when `config` goes out of scope. If the file cannot be saved, too bad - there is no way to signal this without having the dtor throw (which is generally a bad idea in C++).

The member function `node()` return an `s11nlite::node_type` reference, and any serializable data may be put into it or fetched from it.

18.5 micro_api<SerializableType>

This class is one of those, "I'm bored, let's try this out," kind of things. its main intention is to save a small bit of typing (pun unavoidable) when loading or saving the same basic type of Serializable over and over again (as i often do in test code). Here's an example of how to use it:

```
#include <s11n.net/s11n/micro_api.hpp>
...
typedef s11nlite::micro_api<MyType> micro;
micro.save( myobj, 'myfile' );
...
MyType * m = micro.load( 'myfile' );
```

It uses `s11n` to do most of the work, so it inherits options like the default file format. To make the class a tad more useful, it also two other minor features. First, each can use its own file format, set in the ctor or via `micro.serializer_class(classname)`. Secondly, it has simple buffering support:

```
micro.buffer( myobj ); // similar to save() but internally buffered
std::istream is( micro.buffer() );
MyType * m = micro.load( is );
micro.clear_buffer(); // once it's not needed any more
```

19 Memory management and object relationships

"Any day now, any day now, I shall be released."

Bob Dylan

Memory management is an important topic for users of `s11n`. This chapter will try to go into much more detail than i'd really care to about the whens, hows, whys, etc., of memory management in `s11n`. This section is somewhat related to section 16, except that that section covers memory management in the face of exceptions, as opposed to "normal use."

19.1 Data nodes

Data nodes, by convention, are responsible for their own memory management. This means that they own the resources used to store their properties and they own their children. *How* they do that is undefined, but *that* they do it is a given.

For most purposes, data nodes do not need any special memory management. The notable exception is when creating an *unparented* node on the heap (using `new` or `node_traits::create()`). In this case it is often desirable to use a `std::auto_ptr` to hold the pointer until you have a place to reparent it, as in this example:

```
typedef s11n::node_traits<NodeType> NTR;
std::auto_ptr<NodeType> n( NTR::create( 'fred' ) );
... perform some operation which might fail ...
... on success, do: ...
NTR::children(parentnode).push_back(n.release()); // pass ownership to
parentnode
```


19.2 Containers of pointers

Let's consider this simple case:

```
typedef std::list<int *> IList;
IList * il = s11n::load_serializable<IList>( "file1.s11n" );
```

That looks all innocent, but there are some potential pitfalls here. The first, most obvious, is that the caller needs to not only delete `il`, but also the pointers contained in `il`. The library has some utility functions for doing this:

```
s11n::free_list_entries( *il ); // DON'T DO THIS!
delete il; // it's now empty, but...
```

That seems simple enough, but let's look at a subtly more complex case:

```
typedef std::list<IList *> IListList;
IListList * ill = s11n::load_serializable<IListList>( "file2.s11n" );
...
s11n::free_list_entries( *ill ); // deletes all pointers
delete ill;
```

The major error here is, we've leaked the contents of each and every sub-list. We properly deleted the allocated sub-lists, but not their contained parts. A classic memory leak.

This is the main problem with container of pointers vis-a-vis deserialization, *especially* when exceptions are thrown during deserialization. Consider:

```
typedef std::list<MyType *> MyList;
```

During deserialization, maybe the fourth entry in the list fails to deserialize. What do we do here?

Even if deserialization succeeds, someone has to delete those pointers someday. Presumably, this is already accounted for in your application, so the only "danger zone" for these pointers is between the time they are instantiated and the time `s11n` gives them back to your application. In that "danger zone", a misplaced exception could potentially lead to a memory leak.

As of version 1.1.3, the internal exceptions handling was gutted and rewritten to accommodate this type of situation. A "cleanup functor" is now associated with each `Serializable` type (section 6.2.1) to take care of deallocating objects when a deserialization operation fails. The functor is designed such that specializations are put in place to recursively walk any contained sub-parts, so that we can properly clean up even the following type without special client-side action:

```
list< multimap< int, map< string, vector< int *> > > >
```

Clients needing to clean up pointers such a type can do the following:

```
s11n::cleanup_serializable(
myListOfMultiMapOfIntToMapOfStringToVectorOfPointerToInt );
```

Be aware that this is *not* a general-purpose clean-up mechanism: it only works properly if all types involved are registered `Serializers` with proper cleanup functors installed.

When deserializing non-standard containers, you may need to install your own cleanup functors to be sure that entries can be walked and cleaned up if needed.

Some have suggested using smart pointers to eliminate this type of problem, but I don't feel good about imposing a specific smart pointer implementation on `s11n` clients. It is something to consider, nonetheless.

19.3 Cleaning up *before* deserialization

While the core library will never directly do this, it is possible, even sometimes desirable, to do via client-side code:

```
MyType myobj;
deserialize( mynode, myobj );
... use myobj ...
deserialize( anothernode, myobj ); // obtain a new state in old object
```

There is nothing fundamentally wrong with this - it is conceptually identical to a copy/assignment constructor - but there is one immediate implication for authors of deserialization operators: the operators should behave like

copy and assignment operators.

Put simply, deserialization algos must be sure to free up any resources which the deserializing object owns when they take on a new state as a result of deserialization. A common example would be a type which maintains a list of children or values. A simple demonstration of the copy/assignment metaphor:

```
T t1;
... populate t1 ...
T t2;
... populate/use t2 ...
t2 = t1;
```

Assuming "owning copy semantics", at the assignment `t2` would free up any children it currently owns then copy those from `t1`. The same applies to deserialization, which is logically similar to a copy/assignment constructor.

19.4 Cleaning up after failed deserialization

19.4.1 Understanding the problem

It would be nice if we could add text similar to the following in the API docs for every deserialization algorithm:

If this function fails, the target deserializable is not modified and any allocated resources are destroyed.

The problem is, we can't. After going through the code very carefully, trying to figure out where to `try`, where to `catch`, and what to clean up after doing so, it became clear that `s11n`'s architecture blinds it in this regard. Consider this simple call:

```
typedef std::list<T *> TList;
TList list;
deserialize<NodeType,T>( mynode, list );
```

If that fails, we might expect the list deserialization algorithm to be able to clean up any pointers it allocates. This is a reasonable wish, but it cannot be fulfilled. If you read section 19.2, you probably see why, but let's expand on it for a moment:

```
typedef std::list<TList *> TListList;
TListList * tll = deserialize<NodeType,TListList>( mynode );
```

Let's say we have a serialized `TListList` containing 3 `TList` pointers. Deserialization of the first two works, so `tll.size() == 2`. We get to the third one and it throws for some reason. The list deserialization algo can catch that... *but then what?* The natural reaction would be to clean up the whole list of allocated objects. However, if we do that, we end up deleting the `TList` pointers, but not the (`T*`) they contain.

The catch is, deserialization of the `TList` and `TListList` types both go through the exact same algorithm, and the algorithm has no way of directly knowing what it is deserializing - it simply passes the requests to the `s11n` core, which will route them through the algorithms registered for the given types.

This doesn't just affect container types, but any types which hold unmanaged pointers to memory allocated during deserialization. Only the algorithms which work "self-contained", without passing any calls on to other algos or the core, have any chance at all of knowing what they need to clean up on error. Container-related deserialization algorithms must, by their very nature, pass on calls to other algorithms, and therefor cannot normally be self-contained.

The end effect is, they cannot know if they've just failed to deserialized a (`T*`), `list<T*>`, or `map<int,Foo<multimap<double,T*>>>`, and therefore deallocating can never be done safely from that level of the API. Unfortunate, but seemingly unavoidable. The burden of cleaning up on failure then shifts to code which knows about the overall structure of the data (i.e., the client). Or does it ... ?

19.4.2 Accommodating the problem, approach 1 (*don't do this!*)

To extend the above example, let's show where this cleanup needs to be done. In short, the only place which it can be reliably done is from some point which has enough information to know the underlying structure of the deserialized object. In our case, that means a point at which we know about `TListList`. Given that, we might do something like the following in our deserialization operator:

```
try {
    ... deserialize our TListList ...
} catch( ... ) {
    for each myTList in myTListList {
        // free the (T*) in each list
    }
    throw;
}
```

19.4.3 Accommodating the problem, approach 2 (*do this instead!*)

Here is a *much* more general way of managing this problem, at least within the context of Serializables:

```
try {
    ... deserialize our TListList ...
} catch( ... ) {
    s11n::cleanup_serializable( myTListList );
    throw;
}
```

Now we don't care if `myTListList` is a pointer or reference. We also don't care if it's a container or an integer or a `FooManChoo`. As long as the type meets the requirements for the `s11n`'s cleanup functor mechanism, then this will work. The majority of Serializable types need no special support or have that support built in to their registration process. In this specific case, `cleanup_serializable()` will empty out `myTListList` and all sublists, regardless of how many lists or how deeply they are nested, deallocating any pointers in the lists as it goes. See section 6.2.1 for more details.

19.5 Understanding "serialization ownership"

`s11n` was originally designed to enable the serialization of hierarchies of objects. As in any OO design, the relationships of resource ownership are important to concretely define, such that users of the library and the library itself know when each one is in control of a resources (normally this means, "who's going to delete it?"). While `s11n`'s ideas of ownership normally match up nicely to hierarchies of client-defined types, there are cases where users will need to give some thought to questions like:

- For shared resources, who is responsible for de/serializing them?
- How do we de/serialize relationships with shared resources in such a way as to not de/serialize the shared resources multiple times in one transaction?

The general topic of "who is responsible for de/serializing each part" is called "serialization ownership." It is not fundamentally different from normal resource ownership but users must ensure that their de/serialization algorithms' ideas of ownership jive with their internal ownership models, or Grief may show its ugly head. This can range from duplicating objects, leaking some of them, trying to use not-yet-deserialized objects, and so on. So pay attention...

19.5.1 The basic case: objects own their own resources

In many basic OO cases, ownership of a resource belongs to the object which contains it. For example:

```
class Foo {
    SomeT * m_t;
public:
    Foo() : m_t(new SomeT) {}
    ...
};
```

It is fairly obvious that each `Foo` instance owns its own copy of `SomeT`. If we want to de/serialize that member, we have no ownership-related questions, because each `Foo` owns his own `SomeT`. Thus our deserialize operator might look something like this:

```
delete this->m_t; // free up the old one
this->m_t = new SomeT; // create a new one to deser to
s11n::deserialize_subnode<NodeType,SomeT>( srcnode, 'somet', *this->m_t );
```

Or cut out the delete/new and hope that `SomeT` implements careful cleanup when we re-deserialize it.

We could also polymorphically deserialize `m_t` if we need to, by replacing the bottom two lines of that code with:

```
const NodeType * ch = s11n::find_child_by_name( srcnode, 'somem' );
this->m_t = s11n::deserialize<NodeType,SomeT>( *ch );
```

The point is, though, that we own `m_t` and can (should) thus make sure it's clean before deserializing. In this case, our "serialization ownership" is exactly in line with our object's ownership of `m_t`, so we don't have any special concerns here.

19.5.2 Serializing pointers to data we don't own

Let's say we have a class with this private member:

```
list<const SomeT *> m_list;
```

Remember that we cannot directly deserialize containers of `const` objects, as we can't change (deserialize) their states, so that is our first problem. The second problem is, in this case, this object does not own the listed objects, but we still need to serialize our association with them.

This is a trickier case than simple in-object ownership. It can be satisfactorily solved, but necessarily requires some client-side help. Let's outline how we might go about making that list persistent.

In the absolute simplest case, we can deserialize to a `list<SomeT*>` (*non-const*) and then transfer the pointers to wherever we need to immediately afterwards, directly as part of our overall deserialize algo.

In a more complex case, we might need to store a central registry of objects and our relationships to them. Here is one potential way to do that...

First off, we will make some assumptions:

- We have a central registry/pool of pointers to shared objects. Our list contains pointers to those objects.
- The registry associates a unique key with each object and provides an API for searching by key or object pointer.

Certain clients may not need these features, and some may need more. We will start with these, however, to demonstrate a fairly straightforward way of serializing "links" to "external" objects.

When saving our application's state, we will presumably save the shared object pool at the same time. This is fairly trivial to achieve in many cases. Let's assume that our registry internally uses a `std::map<ObjectKeyType,ObjectType*>`, or similar, to store the pool, and that all contained types are `Serializable`. In that case, we can simply use built-in `s11n` support to do what we need:

```
#include <s11n.net/s11n/proxy/pod/int.hpp> // assume ObjectKeyType == int
#include <s11n.net/s11n/proxy/std/map.hpp> // default map proxy
#include 'ObjectType_s11n.hpp' // hypothetical s11n registration
typedef map<ObjectKeyType,ObjectType*> RegistryMap;
RegistryMap map;
... populate map ...
s11n::serialize( targetnode, object_map );
```

Not too difficult.

Now, deserialization of the map inherently keeps our keys associated with the objects, such that deserialization of our downstream objects can find the objects by key (which *they* serialized) later on.

When we serialize our member list, the work is fairly simple (achtung: pseudocode):

```
typedef std::list<ObjectKeyType> KeyList; // string/ulong/etc are likely
KeyList klist;
for each item in m_list {
    klist.push_back( Registry::get_key( item ) );
}
s11n::list::serialize_streamable_list( destnode, klist );
```

Or something along those lines. The idea is, we have a way of looking up some unique key associated with each object, and we simply store a list of those keys.

For deserialization, it's just the opposite, except that now we can populate that `list<T const *>`:

```

this->m_list.clear(); // important to avoid potential extra entries!
KeyList klist;
s11n::list::deserialize_streamable_list( srcnode, klist );
for each item in klist {
    this->m_list.push_back( Registry::get_object( item ) );
    // may be (T const *)
}

```

It is not always that simple, however, as some objects may not be suitable for this type of lookup, or this type of lookup may not exist in your framework, or might be non-trivial (or non-value-adding) to add. In any case, the problem of handling "links" to external data, or de/serialize const data, can often be handled by breaking down the de/serialization into multiple parts. Remember that algorithms can be hidden behind others, so this need not affect the way clients serialize your types, but may affect the internal implementations of the de/ser algos.

19.5.3 Two-way parent/child relationships

A fairly common case for which the above is not a suitable solution is where parent and child objects have an explicit two-way relationship. One common problem here is communicating the parent pointer to a new child during deserialization. This is normally not as problematic as it may initially seem, however, in particular if the parent owns the children pointers. In this case, children do not serialize the link to their parent. Instead, the parent serializes the list of children as normal. During deserialization, the parent does the following:

```

deserialize list of children;
for each child in list {
    child->set_parent( this );
}

```

This of course assumes that the child does not need the parent in order to fully deserialize.

Doing this sort of post-deserialization processing is not at all out of line in using s11n. In many cases it is desirable to manipulate an object directly after deserializing data, in particular when it comes to establishing relationships with objects which were not part of the deserialization operation. For example, while we cannot serialize a network connection, we can serialize the connection parameters, and deserialization could re-establish a connection based on those parameters.

19.6 Known problematic cases

There are a few known caveats to error handling for specific types. Those are detailed here:

19.6.1 std::set<T*> and std::multiset<T*>

In a `std::set<T>`, the contained elements are const. This means that the cleanup process may not modify the items. i.e., it cannot clean them up without violating their constness. Versions 1.2.0-1.2.2 had incorrect cleanup handlers for `set<T>` which would fail to compile in deserialize algorithms. This is "fixed" in 1.2.3 by doing the same thing we do for `std::map` keys: nothing. See below...

19.6.2 std::map<K,V> and std::multimap<K,V>

In a `std::map<K,V>`, the K part is const, so the cleanup routines cannot legally clean it up. Thus the default cleanup handler for these types does nothing for the key part and only cleans up the value part. Since it is unusual to serialize maps with pointer-qualified keys, this isn't seen to be such a big problem.

20 Using plugins

s11n has *rudimentary* support for so-called plugins, which basically means it can load new types at runtime. The primary reason this feature is to allow us to deserialize types which we don't know about at the time an input stream is read. This means that the simple act of deserialization may include arbitrary new types into an application.

As it turns out, the approach used for loading Serializable types dynamically is the same used as loading almost any other type dynamically. This means that the s11n plugins support inherently supports a wide range of uses unrelated to deserialization. This sections is about finding out how to make use of them.

The plugins layer is an optional feature, not part of the core library. The core makes use of the plugins layer if it is there, but can also work without it (but without the ability to load classes from DLLs). The i/o layer can also

make use of the plugins module to load new file handlers on demand.

20.1 Building plugins support

If you are using the supplied build tree, the plugins module is automatically enabled if the configure script finds a DLL loader it can use. On Unix platforms this would be either `libltdl` (preferred) or `libdl` (the *de facto* Unix standard). On Windows, `LoadModule()` is used. If there are problems building it, you can disable it by passing `-without-plugins` to the configure script. See the header files `s11n_config.hpp` and `plugin_config.hpp` for the macros related to configuring plugin support (those files are both generated by the Unix-side configure process, and may need to be hand-edited on Win32 systems).

20.2 Win32 Achtung

The plugins code fundamentally works under Windows, but its usefulness is significantly more limited than under Unix platforms because of Win32's requirement that we explicitly export symbols which we want to be published from a DLL. This means that any types which want to participate in the plugins model must be exported using the appropriate API. See `export.hpp` for the s11n-related macros for this.

The s11n library does not currently (1.1.2) work as a DLL under Windows because of this requirement to export everything.

A related thing to keep in mind is that the classloader model requires that projects building under MS Visual Studio (or similar) will need to turn on the "keep unreferenced code" option in their DLLs, or factory registrations within the DLLs will never happen (meaning the plugins layer won't do anything useful).

20.3 The API

The whole plugin layer is comprised of only one class and 4 free functions in the `s11n::plugin` namespace:

```
class path_finder;
path_finder & path();
string find( const string & name );
string open( const string & name );
string dll_error();
```

The API provide no support for examining the innards of a DLL, only for *finding* and *opening* them. This is because the layer is specifically intended to support classloaders of the type used by the s11n core. Under that model, DLLs publish no specific symbols and we do not keep a handle to them.

Opened DLLs are never closed by s11n, as doing so is *fundamentally dangerous*. When your s11n-using application closes, the OS will free up any DLLs the application opened. This is the *only 100% reliable* way to deal with opening arbitrary DLLs, because the plugin layer cannot reliably know (*nobody* can) which DLL-provided resources are in use when it closes a DLL. (If you're interested in losing a long debate, send me an email arguing that it is possible, *in the generic case*, to know when it is safe to close a DLL.)

To find out if your `libs11n` has plugins support enabled, you can use one of the supplied configuration macros:

```
#include <s11n.net/s11n/s11n_config.hpp>
#ifdef S11N_CONFIG_ENABLE_PLUGINS
# ... do plugin-enabled code ...
#else
# ... non-plugin code ...
#endif
```

20.4 Basic Usage

In fact, there is only "basic usage", not "advanced usage."

Most clients will not need to access the plugin layer directly, but if they wish to, it is intended to be used something like this:

```
#include <s11n.net/s11n/plugin/plugin.hpp>
...
using namespace s11n::plugin;
using namespace std;
string where = open( 'my_dll' );
```



```

if( where.empty() ) {
    cerr << "not found or error: " << dll_error() << cerr;
} else {
    cout << "Found and opened DLL: " << where << cerr;
}

```

If `open()` returns an empty string, one of two things have happened:

1. No such file was found in the search path.
2. The file was found but opening the DLL failed. This normally happens because of incompatible library versions, due to missing dependencies or symbols, or the file is not a DLL at all.

In either case, `dll_error()` should return a descriptive string explaining the problem (it returns the `lib[lt]dl` error string, if possible). The value returned by `dll_error()` is only valid for one call. Per long-standing `libdl` conventions, the internal placeholder for the error message is cleared after this function is called, such that it is guaranteed to return an empty string if `open()` succeeds or if `dll_error()` is called twice without an intervening call to `open()`. On Win32 platforms `dll_error()` returns a string containing the error code returned by `LoadModule()`.

The search path consists of both directories and file suffixes, which may be manipulated like so:

```

path().add_path( '/home/me/lib/mylib/plugins' );
path().add_extension( '.so' );

```

Note that there is nothing about the `path_finder` class which restricts it to being used to find only DLLs. Historically speaking, `path_finder` has often been used as a finder for images, DLLs, and XML files. For example:

```

path_finder p;
p.add_path( '/home/me/.myapp' );
p.add_extension( '.xml:.config:.s11n' );
string configfile = p.find( 'main' );

```

That will return a non-empty string if it finds any of `main.xml`, `main.config`, or `main.s11n`, in that order, in the search path.

Contrariwise, the free functions in the `s11n::plugin` namespace *are* restricted to DLL-related paths and file extensions, by convention.

A default set of library search paths is defined at build-time. Likewise, the file extension for DLLs is set at build-time and depends on your platform. For Win32 it is ".dll" and on Unix platforms it is currently hard-coded to ".so", which is not correct for some Unix-like platforms (e.g., Darwin uses ".dylib"). These settings are defined in `plugin_config.hpp`, and can be modified at runtime using the object returned by `path()`.

21 s11n-related utilities

"I get by with a little help from my friends."

The Beatles

This section lists the utility scripts/applications which come with `s11n`, plus some tools which are known to be useful with `s11n` but are not shipped with it.

21.1 s11nconvert

Sources: `src/client/s11nconvert/main.cpp`

Installed as `PREFIX/bin/s11nconvert`

`s11nconvert` is a command-line tool to convert data files between the various formats `s11n` supports.

Run it with `-?` or `-help` to see the full help.

Sample usages:

Re-serialize `inputfile.s11n` (regardless of format) using the "parens" serializer:

```
s11nconvert -f inputfile.s11n -s parens > outfile.s11n
```

Convert stdin to the "compact" format and save it to outfile, compressing it with bzip2 compression:

```
cat infile | s11nconvert -s compact -o outfile -bz
```

Note that zlib/bzip2 input/output compression are supported for *files*, but not when reading/writing from/to standard input/output⁴². You may, of course, use compatible 3rd-party tools, such as gzip and bzip2, to de/compress your s11n data. Also note that compression is only supported if s11n is built with the optional `zfstream` supplemental library and that library supports the desired compression technique.

21.2 s11nbrowser

s11nbrowser is a Qt-based GUI application for reading arbitrary data saved with s11nlite. It is not shipped as part of s11n, but is distributed as a separate application, available from:

<http://s11n.net/s11nbrowser/>

This tool shows data in a tree structure, allowing the user to browse it. It can also save to all supported formats and can load new formats via DLLs.

22 Miscellaneous features and tricks

"It slices! It dices! It cuts through a tin can as easily as it cuts through a tomato!"

Advertisement for Ginsu(tm) knives

s11n has a number of features which may be useful in specific cases. While some of them require support code from "outside the s11nlite sandbox", a few of them are touched on here.

22.1 Saving non-Serializables

Let's say we've got a small `main()` routine with no support classes, but which uses some lists or maps which we would like to make persistent. No problem - simply use the various free functions available for saving such types (e.g. section 10.4). This can be used, e.g. as a poor-man's config file:

```
typedef std::map<std::string, std::string> ConfigMap;
ConfigMap theConfig;
... populate it ...
// save it:
s11nlite::node_type node;
s11n::map::serialize_streamable_map( node, theConfig );
s11nlite::save_node( node, 'my.config' );
...
// load it:
std::auto_ptr<s11nlite::node_type> node(
    s11nlite::load_node( 'my.config' ) );
if ( ! node.get() ) { ... error ... }
s11n::map::deserialize_streamable_map( *node, theConfig );
// theConfig is now populated
```

Alternately, simply use `s11nlite::node_type` as a primitive config object or the `s11nlite::simple_config` type.

If the Config object is a Serializable object (or a proxied one) it becomes even simpler: simply use the `save/load()` or `de/serialize()` functions directly on the object. For example, to proxy the above map, we could simply insert the following code before we attempt to de/serialize the map:

```
#include <s11n.net/s11n/proxy/std/map.hpp>
#include <s11n.net/s11n/proxy/pod/string.hpp>
// ^^^ map's contained types must be serializable, too
```

In that case, we could use the standard de/serialize functions on the map:

⁴² Sorry, we don't have an in-memory de/compressing streambuffer.

```
s11nlite::save( theConfig, 'my.config' );
...
ConfigMap * m = s11nlite::load_serializable<ConfigMap>( 'my.config' );
if( ! m ) { ... error: file not found or deser failed ... }
theConfig = *m;
delete m;
```

There are other ways to deserialize the ConfigMap object, such as using:

```
std::auto_ptr<s11nlite::node_type> node(
    s11nlite::load_node( 'my.config' ) );
if( ! node.get() ) { ... error ... }
s11nlite::deserialize( *node, theConfig );
```

22.2 Saving application-wide state and Singletons

It is sometimes useful to be able to serialize the state of an application though we have no specific object which holds all application data. This can be handled by defining a simple Serializable which saves and loads all global data via whatever accessors are available for the data. The same approach can be used for Singletons, which we would not normally be able to dynamically load via deserialization due to their very Singleton nature. An example of how to set this up:

```
struct myapp_s11n { // our 'placeholder' Serializable type
    template <typename NodeT> // Serialize operator
    bool operator()( NodeT & node ) const {
        typedef s11n::node_traits<NodeT> TR;
        TR::class_name( node, "myapp_s11n" );
        ... use algos to save app's shared state ...
        return true;
    }
    template <typename NodeT> // Deserialize operator
    bool operator()( const NodeT & node ) {
        ... use algos to restore app's shared state ...
        return true;
    }
};
```

Then register it as a Serializable, which is simpler than for most proxy cases because our "proxy" is actually a Serializable implementing the so-called Default Serializable Interface:

```
#define S11N_TYPE myapp_s11n
#define S11N_TYPE_NAME "myapp_s11n"
#include <s11n.net/s11n/reg_s11n_traits.hpp>
```

To save application state, we simply need:

```
myapp_s11n state;
s11nlite::save( state, 'somefile.s11n' );
```

To load our app state we can take a couple of different approaches, but the most straightforward is probably:

```
myapp_s11n * state =
    s11nlite::load_serializable<myapp_s11n>( 'somefile.s11n' );
if( ! state ) { ... error ... }
delete( state ); // no longer needed - it modified the global state for us.
```

Or, if you want to get fancy, perhaps something like:

```
{ // create a scope to contain an auto_ptr<> object...
    std::auto_ptr<myapp_s11n> ap(
        s11nlite::load_serializable<myapp_s11n>( 'somefile.s11n' )
    );
    if( ! ap.get() ) { ... load failed ... }
}
```

Or, alternately:

```
using namespace s11n_lite;
std::auto_ptr<s11n_lite::node_type> node( load_node( 'somefile.s11n' ) );
if( ! node.get() ) { ... error ... }
myapp_s11n state;
deserialize( *node, state );
```

22.3 Saving lib state plus arbitrary client-specified state

Extending the previous example... i recently had a case which evolved an interesting trick:

My library provides Serializables but no save()/load() functions, because client apps tend to have their own top-level save/load functions. The problem i eventually ran into was that i have a wide variety of unrelated Serializables, and i wanted a common way to save them and my lib state. The reason was simply organizational: my client-side data had dependencies on the lib-side data, and i wanted them to be saved together. This wasn't a problem, per se, but it lead to a lot of code duplicating the same work. The solution was to indeed add load()/save() support at the base-most library level, but do it in a way which allows the clients to bundle arbitrary data with the library data.

Assuming we have a function, `my_lib_data()`, which returns a reference to a library-wide set of data, here's what a lib-level `save()` function might look like:

```
template <typename UserDataT>
bool save( std::ostream & os, const UserDataT & ud ) {
    using namespace s11n_lite;
    node_type n;
    return serialize_subnode( n, "my_lib_data", my_lib_data() )
        && serialize_subnode( n, "client_data", ud )
        && save( n, os );
}
```

And we do the opposite for `load()`:

```
template <typename UserDataT>
bool load( std::istream & is, UserDataT & ud ) {
    using namespace s11n_lite;
    std::auto_ptr<node_type> n( load_node( is ) );
    return n.get()
        && deserialize_subnode( *n, "my_lib_data", my_lib_data() )
        && deserialize_subnode( *n, "client_data", ud );
}
```

Adding overloads which take strings (file names/URLs) is left as an exercise (tip: they can be implemented in as little as two lines each).

22.4 "Casting" Serializables with `s11n_cast()`

Serializable containers of "approximately compatible" types can easily be "cast" to one another, e.g. `list<int>` can be "cast" to a `vector<int>`, or even a `list<int>` to a `vector<double*>`. What exactly constitutes "approximately compatible" essentially boils down to this: the two types must have the same or compatible s11n proxies installed. If the algorithms are written to accommodate it, the pointeriness of the contained types is *irrelevant*.

Assuming we have registered the appropriate types, the following code will convert a list to a vector, as long as the types contained in the list can be converted to the appropriate type:

The *hard* way:

```
s11n_lite::node_type n;
s11n_lite::serialize( n, mylist ); // reminder: might fail
s11n_lite::deserialize( n, myvector ); // reminder: might fail
```

Too much effort. Let's try something *slightly-less-difficult*:

```
s11n_lite::node_type n;
```

```
bool worked = s1nlite::serialize( n, mylist ) && s1nlite::deserialize( n, myvector );
```

Or, the easy way:

```
bool worked = s1nlite::s1n_cast( mylist, myvector );
```

Done!

As of version 1.1.3, `myvector` (the target) is guaranteed to be unmodified if the cast fails.

It is important to remember that only types which use compatible de/serialization algorithms may be `s1n_cast()` to each other. The reason is simply that the de/serialize operators of each type are used for the "casting", and they need to be able to understand each other in order to transfer an object's state. When casting "incompatible" types, `s1n_cast()` will return false or propagate an exception.

22.5 Cloning Serializables

Generic cloning of any Serializable:

```
SerializableT * obj = s1nlite::clone<SerializableT>( someserializable );
```

As you probably guessed, this performs a clone operation based on serialization. The copy is a polymorphic copy insofar as the de/serialization operations provide polymorphic behaviour. To be certain that the proper classloader is used, you should explicitly pass the templated type, using the base-most Serializable type of the hierarchy (the *registered* base type). When cloning *monomorphs* this template typing is not an issue (unless the type may one day become a polymorph, in which case *not* explicitly specifying the template parameter is potentially bug in waiting).

22.6 Half-intrusive proxying and useless friends

This is all theory: i've never tried it, as i don't like C++'s "friend" feature.

It might be tempting to try "half-intrusive" serialization by defining an object which does the serialization, but which has access your type's private data. C++'s friend feature could of course be used to solve this. From the declaration of `MyType`, instead of directly befriending your concrete proxy type, try befriending it via `s1n_traits<MyType>` with:

```
friend class s1n::s1n_traits<MyType>::serialize_funcutor;
```

This ensures that `MyType`'s code doesn't change when his friends do. Sneaky, maybe, but seems reasonable.

There is one small fly in the ointment, though: the de/serialize functor types are, in practice, always the same type, *but are not guaranteed to be*. That means that if we do this:

```
friend class s1n::s1n_traits<MyType>::deserialize_funcutor;
```

Then we are likely to get a warning from the compiler complaining that we've befriended the same type twice.

Note that it is always useless to befriending functions in the `s1n` public API, like `de/serialize()`, because those functions don't actually touch your objects: they only delegate to the types defined in `s1n_traits<MyType>`.

22.7 zlib & bz2lib support

As of 1.1, this support comes in the form of an optional add-on library, `zfstream`, which `s1n` will use if the build process finds it. It can be downloaded from the `s1n.net` downloads page:

<http://s1n.net/download/#zfstream>

When enabled, `s1n` reads `zlib/bz2`-compressed data files without having to know that they are compressed. In the interest of data file portability/reusability, *output file compression is off by default*. Since the feature comes from an external library, the `s1n` API provides no direct way for users to enable compression for output files. It can be enabled client-side by doing the following:

```
#include <s1n.net/s1n/s1n_config.hpp>
```

```
#if s11n_CONFIG_HAVE_ZFSTREAM
# include <s11n.net/zfstream/zfstream.hpp>
#endif

...
#if s11n_CONFIG_HAVE_ZFSTREAM
    zfstream::compression_policy( zfstream::GZipCompression );
#endif
```

Since `s11n::io` uses `zfstream` to open `i/o`streams, `s11nlite` will use the policy specified by `zfstream`.

All functions in `s11n`'s API which deal with input files transparently handle compressed input files if the compressor is supported by the underlying framework, regardless of the policy set in `zfstream::compression_policy()`: see `zfstream::get_istream()` and `get_ostream()` if you'd like your client code to do the same. Note that compression is not supported for arbitrary streams, only for files. Sorry about that - we don't have in-memory de/compressor streambuffer implementations, only file-based ones (if you want to write one, PLEASE DO! :).

As a general rule, `zlib` will compress most `s11n` data approximately 60-90%, and `bzip` often much better, but `bzip` takes 50-100% more time than `zlib` to compress the same data. The speed difference between using `zlib` and no compression is normally negligible, and loading large gzipped files can actually be slightly faster than using no compression. `Bzip`, however, is *noticeably* slower on medium-large data sets.

As a final tip, you can enable output compression pre-`main()`, in case you don't want to muddle your `main()` with it, using something like the following in global/namespace-scope code:

```
static int bogus_placeholder =
    (zfstream::compression_policy( zfstream::GZipCompression ),0);
```

That simply performs the call when the placeholder var is initialized (pre-`main()`).

22.8 Using multiple data formats (Serializers)

It is possible, and easy, to use multiple Serializers, from within in one application. `s11nlite` likes to hide this detail from us, but allows us to set the default Serializer class and load Serializers by class name at runtime.

Traditionally, loading nodes without knowing which data format they are in can be considerably more work than working with a known format. Fortunately, `s11n` handles these gory details for the client: it loads an appropriate file handler based on the content of a file. (Tip: clients can easily plug in their own Serializers: see `s11n/io/serializers.hpp` for the API.)

Saving data to a stream necessarily requires that the user specify a format - that is, client code must explicitly select its desired Serializer. Once again, `s11nlite` abstracts a detail away from the client: it uses a single Serializer by default, so `s11nlite`'s stream-related functions do not ask for this.

Data can always be converted between formats programmatically by using the appropriate Serializer classes, or by using the `s11nconvert` tool (section 21.1).

It is not possible, without lots of work on the client's side, to use multiple data formats in *one data file* - all data files must be processable by a single Serializer. Theoretically, it might be easily achievable if... no, we won't go there.

22.9 Sharing Serializable data via the system clipboard

Experience has shown that holding pointers to objects in the system clipboard can be fatal to an application (at least in Qt: if the object is deleted while the clipboard is looking at it, the clipboard client can easily step on a dangling pointer and die die die). One perhaps-not-immediately-obvious use for `s11n` is for storing serialized objects in the clipboard as text (e.g. XML). Since nodes can be serialized to any stream it is trivial to convert them to strings (via `std::ostringstream`). Likewise, deserialization can be done from an input string (via `std::istringstream`). It is definitely not the most efficient approach to cut/copy/paste, but it has worked very well for us in the QUB project for several years now.

Additionally, QUB uses XML for drag/drop copying so if the drag goes to a different client, the client will have an XML object to deal with. This allows it, for example, to drop its objects onto a KDE desktop.

Assuming you serialize to a common data format (i.e., XML), this approach may make your data available to a wide variety of third-party apps via common copy/paste operations.

The source code for the `s11nbrowser` application contains a class which acts as a global clipboard for `s11n`-able data.

22.10 Containers of const objects

When serializing containers of const objects, we need to do some special-case handling during deserialization. To make a very short example, let's assume that our class contains a list which we would like to serialize:

```
typedef std::list<const MyType *> ListT;
```

That will *serialize* just fine, but *deserialization* will fail at compile-time because the deserialization algorithm of `MyType` expects a non-const object which it may modify. It is an inherent property of *Deserializables* that they may not be const, just as it is an inherent property of *Serializable*s that they must⁴³ be const.

In this case we need to apply the layer-of-indirection rule. One straightforward approach is, in our *deserialize* operator, to deserialize the list to a temporary container of `list<MyType*>`, then copy or move the pointers into your `ListT`, like so:

```
typedef std::list<MyType *> TempT;
TempT tmpList;
if( s11n::deserialize( mynode, tmpList ) ) {
    ... copy/move tmpList's contents to our member list ...
}
```

We must of course be careful with the pointer ownership: `tmpList` owns the pointers initially, and we will need to move that ownership to wherever is appropriate for our application.

Note that it is theoretically possible to add a simple wrapper which handles this const-related handling for a certain class of container (e.g. lists or maps), such that we could do something like:

```
deserialize_list_of_consts( mynode, mylist );
```

The function would need to internally strip out constness from `ListT::value_type`, so it would have some template meta-code, but i believe it could be done with little effort.

22.11 Versioning of s11n data

As discussed (read as "justified") at length elsewhere in this document, i'm not a fan of data versioning. Let's consider one way it might be implemented, and which is fundamentally similar to how the Boost serialization library accomplishes versioning (which it includes in its equivalent of `s11n_traits`):

```
template <typename T>
struct version_checker {
    ... serialize operator uses node_traits::set() to embed a version identifier
    ...
    ... deserialize operator uses node_traits::get() to check the version
    identifier ...
};
```

Now register that type as the proxy for any given *Serializable*:

```
#define S11N_TYPE MyType
#define S11N_TYPE_NAME 'MyType'
#define S11N_SERIALIZE_FUNCTOR version_checker<MyType>
#include <s11n.net/s11n/reg_s11n_traits.hpp>
```

As a final bit, we specialize `version_checker<MyType>` and do any type of validation we like. *Viola*.

There is a caveat, however: you may have to use custom variants of otherwise "standard" *s11n* proxies/algorithms. e.g., the container proxies would not like you adding another property to the target node, and may become angry or confused (throw or result in corrupted node content). To work around this, the version checker could actually restructure the serialized data. For example, our *serialize* operator might embed a new node in the target node, storing the version property in the original target and adding the *serializable* object to a new subnode:

```
bool operator()( s11n_lite::node_type & tgt, const SerializableType & src )
const
{
    typedef s11n_lite::node_traits NTR;
    NTR::set( tgt, 'version', 42 /* need not be an int */ );
```

⁴³ Well, "should" be const. Most serialization libraries do place const requirements on objects being serialized.

```
return s11n::serialize_subnode( tgt, 'data', src );  
}
```

Likewise, the deserialize operator would throw if the version identifier does not match. To avoid duplication of the identifier in both de/serialize algorithms, the identifier might be set as a static const member in the `version_checker` specialization, or made available via a static getter function.

Since this behaviour effectively only works monomorphically, the normal call to `NTR::class_name(tgt, '...')` is unnecessary because it is set by the core.

The remaining caveat involves polymorphic version checking: versioning of types with polymorphic/virtual de/serialization operators effectively requires *those types* to do any version checking themselves, or expose an API which a proxy can use for doing the checks, as the de/serialize implementations otherwise theoretically cannot get at the version info of any *subtype* in the hierarchy.

22.12 Splitting up your output

One of the interesting inherent properties of all Serializables is that they are inherently composable. That is, Serializables can be de/serialized in isolation or within the context of another Serializable. This means that there is no particular reason that we have to clump all of our data into single packets for purposes of saving them. Let's assume that we have a class `AType`, which contains three Serializables, `S1`, `S2`, and `S3`, and that we have public access to the data. The following two approaches are "just as legal" when it comes to saving an object of `AType` to a file:

```
using namespace s11nlite;  
save( myA, 'alldata.s11n' );
```

or:

```
save( myA.s1, 'part1.s11n' );  
save( myA.s2, 'part2.s11n' );  
save( myA.s3, 'part3.s11n' );
```

This is particularly suitable when used with the "saving application state" approach demonstrated in section 22.2.

22.13 Improving compile times

This library's biggest inherent weakness is arguably the compilation-time hit it imposes on client code. Here we will discuss some general guidelines for helping improve compile times...

First include only the proxies which you know you will need. For example, if you're not serializing doubles, don't include a proxy for doubles. For each Serializable we must create a number of back-end types which do things like API forwarding, classloading, etc., using template specializations. Thus the creation of a proxy is not trivial for the compiler.

Secondly, try to reduce your direct dependencies on `s11n.net` headers. Some ways you can do this:

- Create your own front-end interface. With the `s11nlite::client_api` class this is very simple to do. If you know you have a very limited set of types to serialize, for example, all of your classes subclass a base Serializable class, then the native `s11n[lite]` API can be almost completely hidden from client code behind the client-side front-end API. By doing so, we can restrict the `s11n`-related compile times to more isolated parts of our client source tree.
- If you are in the habit of storing declarations in `MyClass.hpp` and implementations in `MyClass.cpp`, then reconsider splitting the implementation file into two files, adding `MyClass_s11n.cpp` (or whatever), and put any parts which contain `s11n` API calls into that file. That way, when the main implementation changes, we don't need to recompile the serialization parts.
- Precompilation might sound tempting, but has a significant inherent flaw: any precompiled source units must be compiled with the exact same options as any client code which will link to it. If this rule is not followed then we run the risk of having inconsistent definitions of anything which might have been conditionally defined/implemented based on preprocessor macros.

22.14 Know when you *don't* need to register a type to serialize it

Members Only. (Most of the time.)

This manual goes on and on about proper registering of types with the framework so it can know how to handle

them. Registrations essentially serve the following purposes:

- Make the classloader aware of the class: its name and how to create an object. This is necessary for polymorphic deserialization. Monomorphic can be done without this.
- Tell the core library which de/ser algorithm pair will act on behalf of the type. Even in the case of types which directly implement a Serializable interface, s11n internally uses a proxy to route itself to the proper class-level API. (This is for internal uniformity.)

As normal in almost all cultures, non-citizens have fewer rights than registered citizens. But they do have some rights. Let's take a look at what they *can* do...

22.14.1 Containers of Streamable types

The following code will work as expected without any registrations of any of the involved types:

```
typedef std::map<int, std::string> Map;
Map m;
... populate it ...
s11n_lite::node_type node;
s11n::map::serialize_streamable_map( node, m );
Map demap;
s11n::map::deserialize_streamable_map( node, demap );
```

The same goes for `s11n::list::[de]serialize_streamable_list()`.

From there we can use `s11n_lite::save()` to send the node to a file, or `s11n_lite::load_node()` to load it from a file.

The reason this works without registration is because the "streamable" algorithms don't need, and don't use, any of the main features provided by the registration process: dynamic loading and mapping of de/serialization algorithms.

22.14.2 Algos which don't need the s11n core API

As a general rule, if we have a type which can be de/serialized *without using features of the s11n core API*, and without dynamic loading, we can get away without registration. We *can* do dynamic loading without the core, but that is an important feature of the library, and there is little reason to want to go around it. By "features of the core," we basically mean any s11n[ite] API which requires a `SerializableType` template argument. The short reason for this is that calling the core library will force us to go through registered proxies (or the default proxy, which won't work in most cases).

In general, the non-registration cases normally exclude any types which have data nested more than one level deep unless we carefully hand-craft out de/ser algorithms to avoid the core API. While it is normally counter-productive to do so, some cases might call for doing this.

A concrete example will help to clarify...

"Streamable" containers, as demonstrated above, work because they explicitly require that all involved types be `iostreamable`. This limitation allows the algos to rely on `i/ostream` operations, rather than the core, to de/serialize each object. Non-streamable containers, however, require registrations for their contained types.

Let's look at why this is so, assuming the exact same map type from the previous section:

```
s11n::map::serialize_map(node, m);
```

There is a fundamental difference between `serialize_map()` and `serialize_streamable_map()`: the former has no idea how to handle the contained types, so it sends them back through `s11n::serialize()`. This, in turn, will attempt to look up the proper handler for the contained type, as defined in `s11n_traits<ContainedType>::serialize_functor`.

Note that if our map's type *is* registered as using the default map proxy, this does the same thing as above, eventually routing through `serialize_map()`:

```
s11n::serialize(node, m);
```

23 Miscellaneous caveats, gotchas, and some things worth knowing

"Don't cross the streams. That would be bad."

Egon, Ghostbusters

23.1 Serializing class templates

Please see the examples on the s11n web site and in the source tree under `src/client/sample/`, which covers this whole process in detail. Fundamentally it is not different from handling any other class, but there are some special considerations which have to be accounted for when registering them.

23.2 Cycles and graphs

While i have *never* seen it happen, it is possible that a cyclic de/serializing object will cause an endless loop in the core, which will almost certainly lead to much Grief and Agony on someone's part (probably yours!). Such a problem is almost certainly indicative of mis-understood or incorrect object ownership in the client code. Consider: presumably only an object's owner should serialize that object, and child objects should generally never have more than one parent or owner.

Data Node-based de/serialization (as opposed to Serializable-based) never inherently infinitely loops because Data Node trees simply don't manage the types of relationships which can lead to cycles. In other words, any such endless loops must be coming from client code, or possibly from client-manipulated Data Node trees.

At least one algorithm has been implemented on top of s11n to serialize containers of a graph of client-side objects, but that particular one was proof-of-concept and it can be implemented *much* better than i have. The point being, it *can* be done, but the library current ships with no algorithms to do this. *If you write one, or even a good, generic description of how to implement one, please submit it!*

23.3 Thread Safety

To be perfectly correct, there are no guarantees. i have no practical experience coding in MT environments, and thus it would be a blatant lie if i made *any* sort of guaranty in this area. But i will tell you what i *think* are the facts...

The s11n code "should" be "fairly" thread-safe, with some notable caveats:

First off, no two threads should ever use the same Serializer instance at the same time: each instance must be used by at most one thread at a time. Violation of that rule is a blanket no-no.

The following Serializers are believed to be 100% thread-unsafe (or un-thread-safe, if you prefer) in all regards:

- `compact_serializer` (reimplementing this one would be quite trivial, but the last thing i wanna do is reimplement yet another damned parser)
- `simplexml_serializer`
- `expat_serializer`

The Serializers `parens`, `funtxt`, and `funxml` have been extensively reworked to use instance-specific internal parsing buffers, as opposed to global data, and are *believed* to be safe in the sense that you may use N instances on N streams from N threads at once. (Let me stress: that is *theory*.)

The guilty code is probably almost all in the flexers, though some of the shared objects (e.g. classloaders) could conceivably be affected. It is believed that the classloader/factory parts, while not specifically thread-safe, are unlikely to be affected by most issues of threadedness. That is, who cares if two threads do a lookup in the classloader at once? The only time this might be a problem is when the optional plugin layer is used, because that layer is akin to `dlopen()/dLError()`, and it is possible that the error string from one thread is read by another.

23.4 Polymorphic types and template parameters

"We've been thinking all these years that Objects and Polymorphism were the solutions to our problems!"

Let's assume we have the following hierarchy of Serializables:

```
T1 <== [extended by] <== T2 <== T3
```

The s11n registration process requires that we register T2 and T3 as subtypes of T1. This is (currently) necessary for proper lookups of the various traitled information, like the proper de/serialization algorithms to use on the type.

Now consider this client-side code:

```
using namespace s11n_lite;
T1 * t1 = new T1;
save( *t1, std::cout ); // fine
delete t1;
t1 = new T3;
save( *t1, std::cout ); // fine
T2 * t2 = new T2;
save( *t2, std::cout ); // oops!
```

The problem with that is that `save()` is going to end up seeing a type of T2, not T1. The end effect is that s11n's core looks to `s11n_traits<T2>` to find out the info it needs, and it may very well not find it. Even if it does, our troubles aren't over: the factory layer probably hasn't got a `factory<T2>` entry, because T2 was registered as a T1 subtype and thus exists in the `factory<T1>`. That means `save()` would work, but loading would not because we couldn't instantiate a new T2 object.

The solution is to template-qualify the call to `save()`:

```
save<T1>( *t2, std::cout ); // fine
```

In practice, this is more of a problem for deserialize/load operations than serialization.

23.5 **Absolute No-nos (Worst Practices) for s11n[lite] client code**

"A muddle of conflicting opinions united by force of propaganda is the worst possible source of control for a powerful technology."

Alan W. Watts, The Book

"It's not a problem until you make it a problem."

Seth Gecko, From Dusk 'Til Dawn

This section, added in version 0.9.17, covers some "no-no's" for the s11n framework. That is, things which are often easy to do but should not be done. They are here because, well, because i've done them more than once and want to spread the word ;).

Please note that the subsection titles below all start with the words *do not* and end with an exclamation point!

23.5.1 **Do not change the name of a passed-in data node!**

`node_traits<>::name(string)` is used to set the name of a node. This name is used by Serializers to, e.g. name XML nodes:

```
<nodename s11n_class='MyClassName'>...</nodename>
```

As a blanket rule:

No code must ever change the name of a node which is passed to it. Code may freely change the names of nodes which it creates.

In any case, when you do change node names, keep in mind that if you want to support the widest variety of data formats, you should follow the standard node naming conventions covered in section 5.3.

An example of this no-no:

```

bool my_algo( s11n_lite::node_type & dest, const my_type & src )
{
    typedef s11n_lite::node_traits NTR;
    // NONO: NTR::name(dest, 'whatever');
    // Never change the name of a node passed to us.
    // The following is Perfectly Acceptable:
    s11n_lite::node_type * child = NTR::create();
    NTR::name(*child, 'foo' );
    // alternately:
    // child = NTR::create('foo');
    NTR::children(dest).push_back(child);
    // or create, name, and reparent in one step:
    // child = & s11n::create_child( dest, 'foo' );
}

```

The reason for not changing the name is essentially this: when building up a tree of nodes, the easiest way to structure nodes (for s11n's purposes) is normally to name them. When a function names a node during *serialization*, the matching *deserialization* algorithm will rightfully expect to be able to find the named node(s). When it cannot find the named node(s), deserialization will likely fail (this depends on the algorithm and data structure, but generally this would indicate a failure). To be perfectly clear: this means that *serialization* is likely to pass by *without error* (in fact, it's almost guaranteed to), but *deserialization* will likely fail (again, "it depends", but it *should* fail).

23.5.2 Do not use a single Data Node for multiple purposes!

See also section 26.2.

Never do something like the following:

```

s11n_lite::serialize( mynode, mylist );
s11n_lite::serialize( mynode, myotherlist );

```

We've just serialized two lists into the same data node (`mynode`). Unless you specifically design algorithms/proxies to handle this, the results are *undefined*. Some algorithms enforce that you give them empty containers, some do not, and the library itself does not specify one behaviour or the other.

Likewise, the following is a related no-no:

```

s11n_lite::node_traits::set( mynode, 'myproperty', myval );
s11n_lite::serialize( mynode, myotherlist );

```

Again, we've used `mynode` for two complete different things: storing a property *and* list contents. If the property is not hosed by the list serialization algorithm then the extra property in the node may very well confuse the deserialization algorithm! Again: *undefined behaviour*. What we need to do in this case is serialize the list into a subnode:

```

s11n_lite::serialize_subnode( mynode, 'child_name', myotherlist );

```

Mixing data from different serialized objects into the same nodes will quite possibly cause a "logical failure" during deserialization. That is, the de/serialization will work, in and of itself, but the results will not be what are semantically expected (but are, indeed, exactly what s11n was told to do). It might work, it might not, depending on a bazillion factors. Don't do it and you won't have to worry about any of these factors.

That leads us to a related no-no...

23.5.3 Do not re-assign a reference returned by s11n::create_child()!

Never re-use a reference returned from `s11n::create_child()` as the target of an assignment to another `create_child()` call. In other words, don't do this:

```

s11n_lite::node_type & n = s11n::create_child( mynode, 'subnode' );
... serialize something to n ...
... Let's re-use n for another subnode ...
n = s11n::create_child( mynode, 'othersubnode' );
// ^^^^ Doh! Just re-assigned the 'subnode' node!

```

That's almost certainly not what's intended. What we probably meant to do was:


```
s11nlite::node_type * n = & s11n::create_child( mynode, 'subnode' );
... serialize something to n ...
n = & s11n::create_child( mynode, 'othersubnode' ); // fine
```

(The changes are marked in blue.)

The design reason that `create_child()` returns a reference is because it returns a non-const which is not owned by the caller (it belongs to the parent node), and i want the interface to intuitively reflect that the caller does not own the returned object. In general C++ practice, object ownership is never transferred to the caller when a function returns a reference.

Another way to create children is like this:

```
std::auto_ptr<s11nlite::node_type> n(
s11nlite::node_traits::create('subnode') );
if( ! (some operation which might fail) ) { return 0; }
s11nlite::node_traits::children(parentnode).push_back( n.release() );
// ^^^ transfer ownership
```

23.5.4 Do not use Serializers to implement classical i/ostream operator functionality!

It may be tempting to implement classical-style i/ostream operators by using `s11n`. The core of `s11n` is i/o ignorant, and using it directly from within your i/o operators is possible, but potentially tedious. The `s11n::io` namespace provides classes which use `s11n`'s conventions to provide a streams-based i/o layer. `s11nlite` provides a binding between the `s11n::io` layer and the core layer. It may be tempting to bypass `s11nlite` and use the `s11n::io` layer from your i/o operators. That is unlikely to work, largely because of the workflow Serializers are designed to follow. Serializers rely on a strict sequence of events which says, "read/write *one* top-level node from/to *this* stream, then you're done." When using Serializers for arbitrary sequences of i/o operators, the Serializer cannot precisely know when a *root* node begins, and thus get confused. If i/o operations are freely mixed in arbitrary order (as they easily could be when dealing with client-side i/ostream operators), the Serializers aren't smart enough to deal with it, as it's far outside of their scope.

Don't forget: if a type is Streamable (i.e., supports i/ostream operators) then it is *inherently* Serializable: if it wants to be treated as a full-fledged Serializable, instead of as a POD, a proxy needs to be installed, such as `s11n::streamable_type_serialization_proxy`. See the various `pod/XXX.hpp` proxy-installation headers for examples of how this is done.

23.5.5 Do not register a type as its own proxy!

Okay, this is not specifically a "do not", but there are good reasons not to do this. Do what? Do this:

```
#define S11N_TYPE MyType
#define S11N_TYPE_NAME 'MyType'
#define S11N_SERIALIZE_FUNCTOR MyType
#include <s11n.net/s11n/reg_s11n_traits.hpp>
```

Proxy objects are created very often - on each call to a de/serialize operator - then immediately destroyed. Unless your type is extremely cheap to create and copy, do not register that type as its own proxy. The default proxies are cheap by design, and have no per-instance state.

Aside from that, this type of registration essentially just doesn't make sense, and no use case to date has shown a need for it. It's really one of those dreaded academic/theoretical problems which is unlikely to ever actually show up. But consider yourself warned, nonetheless.

23.6 Best Practices

This section is not all-encompassing, but touches on some techniques which have proven particular useful (or even necessary) for properly using this library...

23.6.1 std::auto_ptr<T> and s11n::cleanup_ptr<T>

Use these liberally. `S11nNodes` manage all of their own resources, so there are no special considerations for cleaning up nodes. Nonetheless, `auto_ptr<NodeT>` is very often useful, and sometimes even necessary in order to ensure correct behaviour during exceptions.

Many Serializables can be used with `std::auto_ptr<T>`, but `s11n::cleanup_ptr<T>` (section 19) is specialized to clean up Serializables which have special cleanup requirements, like standard containers of pointers.

23.6.2 Know when to use s11nlite and when not to

This manual repeats several times that using s11nlite is easier than using the core library. The “lite” API is, however, intended for high-level client-side use. When writing Serializables, or proxies for them, it is normally a better idea to use the core library, and not s11nlite. This essentially means using the code from the s11n namespace and from “supplemental” namespaces like s11n::list and s11n::map.

That said, there is nothing inherently wrong with a Serializable which hard-codes s11nlite::node_type as its node type. It is, however, often just as easy to use a templated NodeType. This also makes the Serializables compatible with the s11n core API regardless of that NodeType's actual type, as long as it conforms to s11n's S11nNode conventions. Keep in mind that the s11nlite::node_type type is just a typedef, and is not guaranteed to be any particular type except that it will be compatible with the S11nNode conventions.

24 Functional serialization

1.1.3 adds some experimental code for doing some tricks common in functional programming. This is still in its very early stages, but i hope to find some useful functional/metatemplate tricks for adding new features to the library.

While the library generally provides all features which “most clients” need for serialization, there are times when that just isn't enough. While writing custom algorithms is not difficult in and of itself, and normally takes no more effort than a few minutes of time to implement a proxy, it would sometimes be nice to have a simple way to work *within* the library, but *around* its default (or registered/proxied) behaviours. Functional composition allows us to do this by building up functors which themselves encapsulate one or more serialization operations.

24.1 #include ...

Most of the code is declared in:

```
#include <s11n.net/s11n/functional.hpp>
```

24.2 Example: serialize via std::for_each()

As an example, let's serialize a map using `for_each()` and a functor which is applied to each child pair of the map. The “more interesting” parts are [colored blue](#).

```
using namespace s11n;
typedef std::map<int,std::string> MapT;
MapT map;
int at = 0;
map[at++] = "one";
map[at++] = "two";
map[at++] = "three";
s11nlite::node_type node;
```

Given that, we can use functors to call the standard API:

```
ser_f( map )( node );
```

That serializes the map using the default serialize functor (the core s11n `serialize()` function). Its overloaded twin takes a functor argument, so you can specify a compatible algorithm (which means just about any s11n `serialize` algo).

As an example, we can use, e.g., a `for_each()` loop and specify a functor for each child object:

```
std::for_each(
    map.begin(), map.end(),
    ser_to_subnode_f( // functor generator
        node, // target node to place children in
```

```

        "child", // name of each child element
        s11n::map::serialize_streamable_pair_f()
        // ^^ serialize algo, applied to each MAP entry
    )
};
// Now deserialize it using a non-conventional approach:
MapT unmap; // target map to deserialize to
typedef std::pair< MapT::key_type, MapT::mapped_type > NCPair;
// ^^^ kludge: strip the const part of MapT::value_type.first
std::for_each(
    s11n::node_traits::children(node).begin(),
    s11n::node_traits::children(node).end(),
    deser_to_outiter_f<NCPair>( // functor generator
        std::inserter(unmap, unmap.begin()), // output iterator
        s11n::map::deserialize_streamable_pair_f()
        // ^^^^ deserialize algo, applied to each NODE child
    )
);

```

Weird, eh? The weirder part is: none of this requires any s11n registrations of the involved types. But it also doesn't yet work on pointer-qualified types, and registration is currently necessary for that case.

Blabber: Theoretically, some metatemplate tricks can allow s11n to internally distinguish between registered and non-registered types, which may allow the library to handle statically-known pointer-qualified types (e.g., (int), (std::string*), and (MyType*)) non-polymorphically. In English, that means that monomorphs would never strictly need to be registered, whereas currently any non-stack-based allocation requires registration (long story). That's an unproven theory, though. The main problem with not registering is getting a type's name, which we actually ignore in the non-dynamic-load case, anyway.*

The `deser_to_outiter_f()` function returns a functor which sends deserialized objects to an arbitrary output iterator, so it can be used on most containers. For containers which support it, this allows deserializing object to a different order than they are saved in, e.g. by using `std::front_inserter()`. It also allows deserializing from one container type to a fundamentally different type, like `map<K, V>` to `vector<pair<K, V>>`. With the proper binders, we could deserialize from a `map<K, V>` to a `vector<V>`, or *potentially* even a `vector<K>` and `vector<V>` in parallel.

Trivia: the "_f" naming convention was picked up from the Boost.MPL library, and means "functor."

We've also added "_f" variants of all of the major algorithms, like `serialize_f`, `deserialize_f`, `serialize_subnode_f`, etc. These can (mostly) be used directly as proxies when registering a type, one each for the de/serialize functors. In the case of the subnode-based algos, which take three arguments, you need to use a binder functor, like `serialize_to_subnode_f<>`, which essentially converts `serialize_subnode_f` to a binary functor (but see also `serialize_to_subnode_unary_f`).

While s11n has had, since the beginning, the ability to define separate objects as the de/serialize functors, that feature has gone entirely unused until recent experimentation began with functional composition vis-a-vis serialization. If s11n didn't have this feature, all participating functors would *have* to implement both de/serialize operators (as we have conventionally done). There are in fact client-side cases where calling of such functors is ambiguous, which is why the split-functor ability has always been there. Curiously, the core s11n library never has a problem with such ambiguity, and the reason is because it's just forwarding stuff along and the context has already properly strictly defined the constness of all involved objects. In client code this ambiguity cannot always be avoided without another layer of indirection or casting. The point being, having a single functor for each operator turns out to be very useful after all.

24.3 Composing custom algorithms from functors

A slight differentiation on the above approach, we can combine various functors to generate custom algorithms on the fly, as shown below. Assuming we have the same types and objects as shown in the previous example:

```

// define a functor to serialize our map:

```

```

serialize_to_subnode_f<s11n::map::serialize_streamable_map_f>
    algo( "child" );
ser_nullary_f( node, map, algo )(); // Serialize it

// Define deserialization algorithm:
deserialize_from_subnode_f<s11n::map::deserialize_streamable_map_f>
    dealgo( "child" );
MapT demap;
deser_nullary_f( node, demap, dealgo )(); // Deserialize it
s11n_lite::save( demap, std::cout );

```

In the end, `demap` will have the same contents as `map`.

Keep in mind that this is a very trivial example, and work in this area started only in September, 2005. Libraries like Boost.Spirit.Phoenix do some absolutely incredible feats of compile-time composition, and i hope to be able to eventually understand it all well enough to apply it usefully in s11n's API. Functional composition allows us to define our algorithms as inlined expressions, which has interesting uses. One example is that it allows us to serialize the same *one type* using *more than one* algorithm without multiple-registration collisions. s11n's core only allows one registered proxy for each type, and composition allows us a way to bypass the default API marshaling.

24.4 Non-default-constructed proxies

One of the more interesting features which algorithm composition gives us is the ability to use non-default-constructed proxies. We currently have the limitation that proxies are copied, not passed by (const) reference, but this allows at least a minimal amount of at-runtime modification of our proxies.

25 Understanding the costs of deploying s11n

(Why is this section so far down in the manual, when this info really should be up near the top? Because it goes into quite a lot of technical detail which will only be fully understood once the s11n architecture is understood. It's kind of a chicken-egg scenario.)

Having a generic, widely-useful serialization framework at hand means, for me, saving tens to hundreds of hours of work on other project trees. Literally, every time i add s11n support to a project, after 10 minutes of work i can say, "thank gawd that's over!"

But of course all lazy programmers end up paying somewhere... and this section is about the overall deployment costs of using s11n in client-side code. While it may not be conventional for a library to document this type of thing, i feel compelled to tell it like it is, if only to balance out with all the hype i've been spouting about the library up until this point ;).

By "costs" we mean things such as:

- Developer learning time.
- Code refactoring effort (if applicable - s11n support can normally be added to client types *post facto*).
- Compilation times. This is definitely s11n's sorest point, due to its heavy use of templates. s11n versions 1.1+ beat 1.0 hands down in this area.
- Runtime resources: RAM and filesystem space.

To be clear, *all* software has various deployment costs associated with it - this is not a detail which is specific to s11n!

This section will attempt to address these costs, to give potential users of the library a good idea of what they might be getting themselves into... hopefully before they get into it. We will not provide many hard numbers, but we will give an overview of where one can expect to incur at least some notable amount of deployment overhead.

For completeness, we really should compare s11n's costs in at least the following contexts:

- The cost of custom-implementing serialization, as opposed to using s11n. It's safe to say that it is rarely, if ever, trivial to implement i/o support for C++ types.
- Compared to integrating "the average 3rd-party library." This of course varies *widely*, depending on the nature of the lib-client dependency, so a blanket comparison cannot be validly made here.
- Compared to the cost of using an equivalent serialization library.

That last context isn't really fair, because there currently is only one such alternative ;). See

<http://boost.org>, and look for Robert Ramey's serialization library, for the only other C++ serialization framework which currently offers *anywhere near* the levels of flexibility and features offered by s11n. i would guess that Robert's library has similar overall deployment costs as s11n, perhaps even slightly lower, and of course has the advantage of the *massive* peer-review system that all Boost libraries go through. i've tried to objectively compare his library and this one in section 28.

While normally we won't go into specifics of s11n vis-a-vis other alternatives, if only because i only use s11n for all of my serialization needs ;), we will attempt to provide an as-objective-as-possible overview of the general types of deployment costs.

As with *any* software, the cost of deployment is a cost paid almost entirely by the *clients* of that software (who may also be the software's developers, as in the case of "internal" software). i *personally* feel that s11n has a relatively low cost of deployment, particularly when compared to the alternative of hand-coding serialization support into a library. That said, i would be extremely interested in hearing *your* own experiences and opinions (or hard facts!) about s11n's cost of deployment. Suggestions for how to lower *any* aspect of deployment costs are always welcomed. :)

25.1 Learning curve

It would not really be fair for me to comment on this aspect of s11n. As its author, i inherently know how s11n works and how to use it, how *not* use it, and how to abuse it. But i *will* of course comment on it, otherwise this section would end immediately after this paragraph.

It is my belief that experienced coders who start with the sample code in the s11n source tree and browse through the docs can pick up the library, almost to the point of full proficiency, within a day or two (maybe faster, for you especially clever ones out there). It can be understood to the point where one can basically use it in a couple of hours or less, i would think. (*If i am way off here, please let me know!*)

My "experienced guesstimate" would say that coders who have posted to the s11n mailing list normally seem to feel comfortable with the architecture after writing 2-3 serializable implementations or serialization algorithms. i can't say how physically long that maps to for beginners - an experienced s11ner can crank out such an implementation for a typical client-side type in a few minutes in most cases.

Please, please, *please*, if you are just starting out with s11n, *start with the s11n-lite API!* See section 2.5 for why.

True masterhood of the library can take time, but how much is unknown and probably unknowable. i will admit that i do not yet fully comprehend all of the *potential* uses, abuses, and tricks implied by the architecture.

There's still a lot of room for theory in there, and at least as much room for experimentation. It will be a while before s11n's current model is worn out, i think (i hope!). Exploring those aspects is half of the fun of working on s11n.

There *is* a *lot* of documentation for the library, but that is *not* because it's hard to use. That is, rather, because:

1. As a client-side software user, i refuse to use undocumented libraries, with a strong preference towards *well-documented* libraries (e.g. Qt (<http://www.trolltech.com>) is a great example, as are the libraries available from <http://boost.org>). Being so pedantic on this point, i cannot expect users of my software to give it a second glance if it's not documented, and not to give it a third glance if simple things like pointer ownership aren't documented. You wouldn't believe how much software does not document pointer ownership. Aaarrggg.
2. Experience shows that documenting software helps to find weaknesses in the API. e.g. if something is difficult to document clearly, it's almost certainly difficult to use properly. Holes in the API have often been caught by documenting the related APIs.
3. i enjoy writing about topics which interest me, and s11n obviously interests me.

Users are *not* expected to read the full documentation in order to be able to use the library, but it is hoped that the documentation will be able to answer most or all of their questions, should they need a reference. If the docs don't suffice, feel free to email us your questions (the address is at the top of this document).

25.2 Intrusivity (or not)

"I hate writing apps around technologies like CORBA and Oracle(tm) [database system] because they force the developer to focus so much on the specifics of that technology, instead of on solving the problem at hand."

Anonymous Software Developer

s11n goes to great pains in order to be as non-intrusive as practical on client code. Clients wishing to support a "conventional" serialization API, where classes derive from some Serializable base type, will of course require

some level of hard dependency on s11n. Clients who use s11n's proxy support can, in many cases, add serialization without having to change their core project code *at all* - rather, they simply need to register the appropriate proxies. Using the proxy approach can help keep client-side dependencies on s11n down to a handful of places, and allows clients to ship s11n support for their classes as an optional component.

25.3 Compilation costs

Yes, i *actually do* have something very negative to say about libs11n: client-side compile times absolutely suck. This was especially true in versions before the mid-0.9 series, and is still a sore point for 1.0.x. It has been improved significantly in 1.1. A simple benchmark program is in the 1.1.3+ source tree:

`src/client/sample/compspeed.cpp`, and the source file includes the results from my PC.

The reasons for the horrible compilation times boil down to:

- We internally create many small template types during compilation to achieve "compile-time polymorphism" and factory registrations for the s11n API. The former is required for API marshaling, amongst other things, and the latter is required for dynamic creation of objects during deserialization⁴⁴.
- Compiling template code inherently takes more compiler horsepower than non-template code, especially when advanced features like partial template specialization are used.
- Compiled template code inherently generates much larger object code than non-template code does. This means longer link times, to resolve multiple copies of templates. This also means *significantly* larger object files, which inherently means more i/o is required by both the compiler and linker(s). Whether or not this specific aspect plays a significant build-time role is arguable, and has never been benchmarked, but it is at least worth mentioning and cannot be completely ruled out as a problem point.

In the 1.0 tree, the main culprits for chewing up compile times are the various proxy registrations: it goes overboard and installs many of them in cases where it doesn't need to in order to simplify client-side usage. In the 1.1 tree we have factored out the proxy registrations into as small of units as are practical. This requires a bit more forethought on the developer's part, as he must decide which headers/proxies he needs to include, but the compile-time benefits should be noticeable in the vast majority of client-side cases. At least, it is hoped that they will be more tolerable :/.

Again, my apologies for the slow compiles, but i simply don't see a way around this problem without doing things like build-time code generation, where we could build the s11n-related code one time in a separate module. Code generators are out of the question, as far the s11n core goes, because they are not in-language. That said, clients are free to do whatever code generation they feel they need to. By pre-generating s11n proxies and compiling ALL s11n support into localized object files, it is theoretically possible to shift the compile-time hits to only those modules. Theory, that is, complicated by the nature of template instantiation rules. If you pull it off, please share with us how you did it.

The book *C++ Template Metaprogramming* [CTM2005] gives some real-world comparisons of compile-time costs of deploying template-based code. While i do beg to differ with some of their numbers (which don't show any significant slowdown until hundreds of types are used, which is much at odds with what i daily see in s11n), it is the only relatively full-fledged analysis i've seen on this aspect of template-based code.

25.4 Memory/RAM costs

Here we will focus on the *theoretical* and *abstract* costs of system memory (RAM) *vis-a-vis* serialization via s11n. Filesystem space is not a special concern in the context of s11n, as filesystem limits apply to any code which saves data. That said, s11n's i/o layer does no unusual tricks, using only the standard iostreams interfaces, so s11n should not exhibit any sort of "unusual" file access costs. Likewise, it does no unusual memory-related tricks like reimplementing `new` or `delete`, or using custom allocators.

At an abstract level, serializing an object requires that we make a logical copy of the object. This is of course not cheap, even if only because Serializable objects have, by their very nature, some number of data members. In abstract terms, let's naively assume that the copy is twice as large as the original. In concrete terms, this is highly unlikely to be the case: the serialized data of course has its own internal overhead. To understand what this overhead might look like, let's take a look at one possible implementation for an S11nNode type, keeping in mind the basic requirements placed on such types by s11n (section 4.2). A basic implementation, not optimized via reference counting, etc., may very well contain the following private data members:

- Two `std::strings`: one to hold the node's name and one to hold its logical class name.
- One `std::list<NodeType*>`, or similar, to hold the children of the node.

⁴⁴ That's not entirely true as a blanket rule for deserialization, but it is a rule for s11n's implementation. We could ditch the factory layer if we either had no, or very limited, support for polymorphism. That's not acceptable, of course.

- One `std::map<string, string>`, or similar, to hold the key/value pairs of the node. Remember that s11n internally uses lexical casting for POD-type type conversion, so internally all properties are stored as strings. While this might *sound* horrible, this is a simple fact of life and also exists in the world of XML, so i don't feel one bit bad about it. (Besides, most `std::string` implementations are optimized a *lot* better than most people give them credit for.)

When serializing lots of small objects, this *might be* huge amount of overhead, relatively speaking. i explicitly say "*might be*" because it *really* depends on factors like reference counting, etc., in your STL implementation. As far as i am aware, all STL implementations use such features in their `std::string` classes. Since s11n uses strings *extensively* for storing raw data, s11n can indirectly benefit from such features if your STL provides them. In any case, as the size of the Serializable object goes up, the relative memory overhead of serializing many of them drops. This is little consolation, i understand.

In addition to the memory cost of strings, there is the runtime cost of lexical casting. For string-typed properties a lexical cast is a no-op⁴⁵, but properties are often not *natively* stored as strings. e.g. in `MyObject`, we might store the `change_time` property as a `long int`, and de/serializing that property will cause a short detour through an ostream operation (for serialization) or istream op (for deserialization).

To be clear about all of this "massive overhead", though, consider the following client-side call:

```
s11nlite::save( myobject, std::cout );
```

Before that function is called, and after it returns, the notorious "second copy" does not exist in memory: it only exists for the life of the serialize operation, and it is thrown away like a used tissue before that operation returns. That is: the cost is an s11n-internal one, and of no *direct* interest to the user, but the user should be aware that serialization *will* eat up memory proportional to the size of the objects being de/serialized (what *exactly* that proportion is, is probably unknowable for all practical purposes).

Remember, too, that client-side objects often also have internal data which is *not* serialized, so the idea that a serialized copy is heavier than the original object certainly does not apply in all cases (mainly it applies to small types - those with only a few POD data members or one container).

Deserialization *normally* has similar costs: we must build up a tree of nodes and populate an object with the data (creating the object if needed). Where there might be a big difference is the specific i/o handler: if it buffers all of its input before it begins deserialization then the memory costs jumps, theoretically/abstractly by approximately another factor of roughly 1x. That is, it is potentially possible that a deserialization results in effectively 3x the memory of an object (again, *very roughly* guesstimated). In practice this 3x explosion should be extremely rare or non-existent because:

1. All of the shipped serializers do no special input buffering: they read input stream-wise, creating nodes as they go, until EOF or they load one complete root node. This is "buffering" in the sense that we transform the stream content to s11n nodes before passing it back into the framework for deserialization proper, but we do not keep the stream content - it is discarded directly after consumption.
2. In deserialization we either have an object to deserialize directly into or we have to create one. In either case we have the same as with serialization: effectively two copies of the object's data. The only difference is that in the dynamic-load case we first build up the node tree and then the object, which is of course the opposite of serialization.

There are cases, e.g. networking, where buffering a whole object tree in a string might be required or might otherwise greatly simplify other code. The `s11nlite::micro_api` class provides a simple API for buffering Serializables to and from strings.

It would be interesting to explore a "destructive" i/o API, in which:

- During serialization, we destroy each node directly after sending it down the i/o pipe.
- During deserialization, we destroy the node directly after we deserialize its contents.

These operations are not possible with the current API due to the required constness of various data. Such operations might also require either new de/ser algorithms or new conventions to accommodate, e.g. a post de/ser functor which algos are required to call on each node. In any case, at some point during serialization we would have a full second copy, but only for a fraction of the time (while de/serializing the deepest leaves of the object tree, since we dive in depth-first). If i/o support were added directly to a `Data Node` type and we add such a "destructive" API, then it might be possible to completely eliminate all second copies, at least at the root level of an object tree (we might need copies of individual objects). Such support, however, is considered project-specific, and well outside the bounds of the core s11n API. That said, the general s11n model *might be*

⁴⁵ In API terms s11n doesn't know the difference between `string` and `int` and `AlaskanPolarBear::MatingInfo`, but some internal optimizing is done to ensure that strings go through as little translation as possible. All that happens, in a worst case, is a `std::string` copy, which is known to be reference-counted in most (all?) STL implementations.

amendable to such an option, perhaps with a little hacking.

25.5 Runtime speed: s11n and the "Big O Notation"

It is architecturally not possible/practical/feasible to impose maximum runtime requirements on the s11n API. For example, we cannot impose the blanket rule that all serialization algorithms must perform their duties in any less than linear time. Stream i/o is one of the places where we simply won't be able to get around paying *at least* linear runtime costs. Client-side algorithms are free to do whatever they like to improve performance, but linear is theoretically as fast as it can get.

As a general rule, most de/serialization algorithms inherently have effectively linear complexity with some constant overhead, but as they may call arbitrary de/serialization algorithms in the course of recursive serialization, they can make no guarantees in this regard. One known exception to the "linear guideline" is the Serializers which do entity translation on their property data (most do this to some degree). The "generic" entity translation algorithm use by s11n is known to perform slowly. i can't name an O notation for it, but it's not a pretty one in any case. i would be extremely happy if someone would contribute a more efficient implementation of `s11n::io::strtool::translate_entities()` :).

i will openly admit to having never comprehensively benchmarked nor profiled libs11n. i have run some small speed tests on my standard 1.4GHz PC, and the numbers were well within what i personally consider to be reasonable. For example, an average load-from-stream rate of 20k-50k object nodes per second, depending on the Serializer, and saving is normally faster. Paul Balomiri, an Austrian s11n user, reports using s11n for some 10 million data nodes, 1 gig of XML data, taking 3 minutes to load: this works out to 55knodes/second, which is close to my numbers (but far, *far* larger than my data sets).

In my opinion, the fact that Paul can get 10 million data nodes in memory at once without thrashing his system to death really says something about his STL implementation, considering the theoretical memory cost of each node (as explained above). i ashamedly admit that i was shocked and happily surprised at finding out that s11n survived Paul's data set.

i personally use s11n in over half-a-dozen projects, none of which have *nearly* the data requirements of Paul's project. i typically save lists and maps, often nested 3 or 4 levels deep, and very rarely more than 10-20k objects (and normally less than a few hundred). Again, i haven't benchmarked save/load times, but "to my eyes" s11n appears to be fast enough to suit the vast majority of client needs. In any case, i cannot say that i have ever felt that the load/save times are "too long" - they seem well with reason to me, from a user's point of view.

That said...

There are ways to help speed up s11n if you are willing to look into options like using a customized S11n Node type or implementing your own Serializer interface (or subclass). The core library is quite small and 99.9% template code, so it may benefit from compiler optimizations. The internals of an S11n Node could be implemented more efficiently if one is familiar with that level of optimization (i'm not, really), and the i/o-related code could certainly benefit from some optimization as well. Keep in mind that s11n's core does not rely on the `s11n::io` code in any way, but that s11nlite does. This means that you can use the provided core and your own i/o interfaces if you like. Users who think that such i/o or node type customizations might be interesting options to explore should feel free to get in touch with us through the development list and we can discuss some potential options.

25.6 Code maintenance costs

"Code maintenance", in this context, essentially means, "how much time one must write s11n-related code." All software has maintenance costs, and these costs are not always trivial.

It is my firm belief that making s11n any less costly, in terms of maintenance, would be extremely difficult to achieve. In the half-dozen or more projects i currently use s11n in, the s11n-related code is effectively write-and-forget. Once an object is Serializable, it's always a Serializable, and is usable in all s11n contexts using the same APIs as all other Serializables. Thus once that code is in place and known to work, it normally becomes a pure background detail.

With the same major-minor number of s11n, major conventions will never be changed, so there shouldn't be significant maintenance-related costs in upgrading. Within a development tree, or between, say 1.0 and 1.2, then 1.2 to 1.4, nearly anything might change, so upgrading s11n might have porting costs.

Changes as major as an architectural overhaul would be denoted by changes in the major number. In that case, of course, there may be any amount of porting costs.

25.7 Money

It would be naive to say that deploying s11n is free of monetary costs. As the old saying goes, "time is money", and thus the general rule is:

s11n's monetary cost of deployment is equal to *your* hourly cost of software development.

That is, every minute of your time it takes you to deploy s11n costs you (or your clients, or someone) one minute of time. Whether or not that time actually costs anyone money or not is not the point - the point is that deploying anything costs someone some amount of their own personal time slice. (Now if i only had 50 cents for every hour i've spent working on s11n...)

The time-is-money equation is of course nothing new, and applies to *any* software deployed *anywhere*. But we're not here to discuss just any software, are we?

i personally consider s11n to have a lower-than-average deployment cost than most Open Source libraries. The main reason is touched on in the previous section: most client-side code is write-and-forget, rather than write-and-maintain. This means, for example, that implementing a serialization algorithm for a given type (or family of types) is a one-time effort. The exact time it takes to write such an algorithm depends on the complexity of the problem, of course, but by taking advantage of existing algorithms for commonly-understood structures, like the STL containers, we can cut coding times even further. For example, proxying *and* saving a `std::map<int, std::string>` equates to approximately the following code:

```
#include <s11n.net/s11n/s11n-lite.hpp>
#include <s11n.net/s11n/proxy/std/map.hpp>
#include <s11n.net/s11n/proxy/pod/int.hpp>
#include <s11n.net/s11n/proxy/pod/string.hpp>
s11n-lite::save( mymap, std::cout )
```

So, the overall money cost can be answered with this question: how long does it take you to do those steps?

As far as the effort it takes to make the average class Serializable - i normally need 5-15 minutes to include all the proper headers, register any proxies i need, write the code, and do basic tests. Registering proxies for well-understood types - e.g. the standard containers (again) - is a job of under 2 minutes, even when typed by hand from scratch. Again, once these registrations are in place, they are background details which needn't worry anyone any more. Granted, i know the library intricately, but from my client code i behave as client code should (that is, exactly what documentation says to do), and thus in principal any experienced coder can churn out s11n algorithms quickly, and therefore cheaply, once they have done it a few times.

26 Common problems

"I pre-emptively accept that from some perspective, these absolutely suck."

Rob Donoghue

In this section i impart some of my hard-earned knowledge with the hope that it saves some grey hairs in other developers...

26.1 Satan speaks through the console during compilation

If, during compilation, your terminal is filled with what appear to be endless screens of gibberish from the mouth of Satan himself, don't panic: that's the STL's way of telling you it is *pissed off*.

It may very well be one of these common mistakes (i do them all the time, if it's any consolation):

- You're trying to serialize a type which isn't yet registered with s11n. This often happens when serializing containers: remember that the contained type(s) must be Serializables, and that a map's value_type (a pair type) must also be made Serializable in order to make a map Serializable. This will normally show up as an error saying that no operator()(something) is defined for the type.
- You've swapped the arguments for a de/serialize() call. By convention, nodes always come before Serializables in the parameter list. Swapping these will cause you no end of error messages from Hell, with things like, "no such list<..>::impl_class()..." or "list<..>::children()". The first hint that the args are swapped is that it's trying to call a `node_traits` function on your Serializable.
- You've tried to pass a pointer as a node argument. *Serializables* are accepted by the `core serialize()` regardless of whether they are passed as pointers or not, but *nodes* are only passed by reference. Why? Because nodes are easy for the API to control in this regard and Serializables aren't, so Serializables get some extra leeway (besides, it was trivial to implement the pointer-to-

reference translation in SAM). This property internally simplifies many operations on Serializables, as well.

- You're trying to pass a (Serializable *) to an algo which does not want a pointer, and this is showing up as a failure in the ability to convert between (Serializable*) and (Serializable). Double-check your calls to algorithms other than the core `serialize()` algo. As of 1.1.3, there is also a `deserialize()` variant which accepts a reference to a pointer.
- You have jumped from `s11nlite` to `s11n` without being aware of the different template args required by like-named functions in the `s11n` namespace. Shame on you. Almost without exception, the `s11nlite::` functions with the same name as `s11n::` functions are missing one template parameter (the first one) - the data node type - because `s11nlite` hides that abstraction. That said, in many cases the calls are identical, because template type resolution will do the right thing, in which case the `s11n/lite` functions are basically the same. `s11nlite` duplicates/forwards lots of functions simply to keep a whole usable client-side API in that namespace. Be sure to check for differences before freely switching between the two (see the API docs).
- Const errors during a de/serialize call: make sure that your Serializable's [proxy's] serialization operators have the proper constness, as defined in section 5. In the case of a proxy, you may have to split it into two functors: one each for de/serialization, and be sure to add `#define S11N_DESERIALIZE_FUNCTOR ...` to the registration call. This should rarely, if ever, be absolutely necessary, however.
- When fetching a child node during a *deserialize* operation using, e.g. `s11n::find_child_by_name()`, be sure you use a (const NodeType *) and not a non-const (NodeType *), as the parent object is const in that context.
- When iterating over containers, be sure to use `const_iterators` if the NodeType or SerializableType passed to the function are const, as appropriate.

To be honest, though, those are just the common ones - any minor violation in usage will cause the STL to go haywire, as I'm certain you have already experienced many times in your coding life. The important thing is to remain calm and simply try to understand what the compiler is telling you. Often a single STL usage error can lead to literally *tens of kilobytes* of error text (I was once punished with 70k for making a *one-letter typo*), but after eliminating the first error the others are likely to go away. Elimination of the problem is normally straightforward once the STL-speak is decoded.

26.2 Containers serialize, but fail to deserialize

See also section 23.5.2.

This is almost invariably caused by a simple logic error:

(Been there, done that.)

When serializing containers, it is essential that each container is serialized into a separate node. After all, each container is ONE object, and one node represents one object. It is easy to accidentally serialize, e.g. both a `list<int>` and `map<string, string>` into the same node, but the result of doing so is undefined. That is, it will serialize, but deserialization may or may not work (don't count on it!).

If you've done that, there may be two ways to recover from it (assuming you need to recover the data):

- Edit the output file and split the nodes up manually. The feasibility for hand-editing depends on the Serializer used: some are not hand-editable. Tips: `s11nconvert` (section 21.1) can convert it to other formats and `s11nbrowser`'s cut/paste features might be useful here (section 21.1).
- Programmatically fish the data out of the node, e.g. using `s11n::find_children_by_name()` to separate the various children. In a worst-case (all entries have the same name, or names are nondeterministic) you'll need to do it based on `node_traits<>::class_name()`, but that would be no fun at all, as they are unpredictable. (Expecting an "AType" node? Think again - you got a "BType"!)

Also, it is essential that you always use complementary de/serialization algorithms/proxies. For example, if you use `serialize_streamable_map()` to save a map, then use ONLY `deserialize_streamable_map()` to deserialize it, as any other algorithm may structure the serialized data however it likes, as defined in its documentation. Be aware of each algorithm's weaknesses and strengths before settling on it, because changing later may not be feasible (old data won't be readable without, e.g. special-case code to check for it and use the "old" algorithm - but such compatibility checks are possible using `s11n`'s proxying model).

26.3 Abstract Interface Types for Serializables

`s11n`'s classloader can handle abstract Interface Types: simply add this line before including the registration code:

```
#define S11N_ABSTRACT_BASE
```

That's all. This does not have to be added for subclasses of that type.

For the curious: this installs a no-op object factory for the type, as those types cannot be instantiated, and thus cannot be created using `new()`. As far as the classloader is concerned, trying to instantiate an abstract type simply causes 0 to be returned.

26.4 Statically linking against s11n

First an over-simplified review of C++ linking rules, then an explanation of the problem:

- When linking object files in with your project, the linker binds all symbols in the object file to your application/library.
- When dynamically loading (opening a DLL), the dynamic linker typically resolves all symbols in the DLL at once (on Unix systems this is normally configurable by passing an additional parameter to the DLL opener, but for most purposes, complete symbol resolution at open-time is the safest bet).
- When linking a static library to your app/lib, the linker only links in symbols which your app/lib actually references (even indirectly).

It is this last point which is problematic for us. This library's classloader model, which is described at length in other documentation, requires that the so-called *static initialization phase* happens in order to instantiate some simple objects. These objects do nothing more than register other types with a classloader, and then they are destroyed and never used again. They are never referenced, per se, and thus they are effectively non-existent when built in to a static library.

When linking statically against s11n, the factory registrations used by, e.g. s11n-lite, will never happen. That means s11n-lite won't know about any Serializers, which means you won't have any i/o support.

A workaround is to link in the object files from the s11n source tree, instead of linking against libs11n's DLL.

27 Evangelism

"If I can sell tickets to Red Sonja and The Last Action Hero, I can sell almost anything."

Arnold Schwarzenegger, while running for governor of California

<http://news.bbc.co.uk/1/hi/business/3374805.stm>

"I want to make sure [a user] can't get through ... an online experience without hitting a Microsoft ad."

Steve Ballmer (Google for one of many URLs)

Obviously, i've got a lot to say about s11n. i mean, how many other Open Source projects of this size have complete API docs, a web site full of example code, *and* a manual of this size ;).

So far i've tried to keep the hype down, but it's sometimes difficult :). In this section i will let loose and explain, in no particular order, some of the library's features which i find particularly interesting, useful, or just downright cool.

27.1 Pointer/reference transparency for Serializables in the core API

That is, the following are equivalent, assuming `list` is a pointer type:

```
s11n::serialize( mynode, list );  
s11n::serialize( mynode, *list );
```

One s11n contributor, martin krafft, is always trying to talk me out of this, but the fact is, that subtle feature allows some really amazing code reduction benefits elsewhere. For example, consider what we would have to do for proxies if they had to expect either a pointer or a reference to a Serializable? You got it: we'd have to duplicate every serialization operator for every serialization proxy. No chance i'm gonna tolerate that, so the pointer/reference transparency stays. It is implemented, by the way, via a single template specialization for SAM (a few lines of code). The reality is that these few lines of code *greatly* reduce maintenance costs

elsewhere. See the map/list algos, all of which handle pointer and value types with the same code, for some examples of what this allows us to do. Or just read on to the next section, where we evangelize just exactly this technique...

27.2 Container-based algos which are pointer/reference-neutral

Consider these two data types:

```
typedef list<string> StringList;
typedef list<string *> PStringList;
```

i banded by head for quite some time to try to figure out how to do de/ser those via one algorithm. That's not as straightforward as it sounds because for deserialization we need to dynamically load the pointer types, and do so polymorphically when possible. Type-dependent branching isn't always *syntactically* possible in C++, so the proverbial *another layer of indirection* was needed to solve the problem of "unified code" for pointers and references. Since the CL layer did the dynamic loading, i wrote up some templates to hide the syntactic and de/allocation differences between pointer and reference types, sticking the CL part behind the pointer-based branch and essentially doing nothing in the reference branch⁴⁶.

After some effort and experimentation, a single pair of remarkably small algorithms evolved, and they now take care of de/serializing any standard list, vector, and multi/set. That is, the following operations all go through the exact same few lines of code to do their work:

```
StringList * slist = new StringList;
PStringList * plist = new PStringList;
// ... populate lists...
s11nlite::save( slist, std::cout );
s11nlite::save( plist, std::cout );
s11nlite::save( *slist, std::cout );
s11nlite::save( *plist, std::cout );
```

That demonstrates two separate s11n features: core API transparency for pointers/refs to `slist` and `plist`, as covered above, and algorithm-level pointer/ref transparency for the `(string)` and `(string*)` elements of the lists. The function `s11n::list::serialize_list()` currently does *all* list-based serialization for the framework (that's a LOT). Likewise, `s11n::list::deserialize_list()` does *all* of the deserialization. (Reminder, that's the *default* implementation, and it can be replaced for any specific container type.)

Not impressed, eh? Let's look only at lines of implementation vs. functional scope:

- `serialize_list()` is implemented in approximately 11 lines of non-debug code.
- `deserialize_list()` has approximately 20.

Consider type L, which is any type conforming to the most basic `std::list` conventions (this also covers `vector`, `deque`, `set` and `multiset`). Now consider the type ST, which may be any `Serializable` Type, *including* L. With the above algos we may generically de/serializer any combination of:

```
L<ST>
L<ST*>
L<L<ST>>
L<L<ST*>>
L<L<ST*> *>
L2<L<L3<L4<ST*>>>
```

ad infinitum...

Get the point?

Now consider that we can do the same, using exactly two algorithms, for any combination of standard map-style types (out of the box that's `std::map` and `multimap`, but client-side map-likes can also work with these algos). Let's assume M is a `map[SK,SV]`, where SK and SV are both `Serializable` types. Now let's begin to look at that more closely, mixed with the `Serializable` list type (L) from the above examples:

```
M<SK, L<SV>>
M<SK, SV>
M<SK *, SV *>
M<L<SV>, L<M<SK*, SV>>>
```

⁴⁶ That "nothing" turned into a long-standing bug-in-waiting, reported by Patrick Lin, which was fixed by adding a one-line "something" in 0.9.17.

ad infinitum, ad nauseum...

and Amen, brothers!⁴⁷

By including the proper proxies, client code gets immediate access to all of the above combinations, plus the *trillions* more they imply. Clients *do* pay compile- and link-time costs, plus fatter binaries, to be sure, but the ease-of-use and coder-effort benefits are, in my opinion, difficult to improve upon. Hopefully, future compilers or development techniques will allow us to cut the compile-side costs. And if not... we'll just need faster PCs ;).

Please note that i'm *not* touting the cleverness of the algorithms themselves, but the flexibility of the s11n architecture, which allows such generic algorithms to plug right in.

If the dimensions of the possibilities don't seem *cool* to you, then s11n probably can't impress you at all (which is all fine and good, i mean - to each his own opinion). However, since this is the Evangelism chapter, i'll go ahead and say: it is my firm belief that s11n supports, *out of the box*, more combinations of data types than most serialization frameworks *could ever hope* to be able to support *at all* (and even then only with unrealistic amounts of client-side or support code). The main reason for this is that s11n takes blatant advantage of newer C++ features which many mainstream libraries shy away from, often for compiler portability reasons. My take on compiler portability is simply this: if we want to save 21st-century data types effectively and flexibly, we need to start using 21st-century tools and methodologies. :-P

27.3 "Casting" between "similar" types

Due largely to the above-mentioned features of pointer/reference transparency, s11n allows us to convert to and from "similar" types with ease (though not necessarily with great efficiency). Witness:

```
list<SomeT *> dlist; // SomeT is any Serializable
vector<SomeT> ivec;
// ... populate ivec ...
assert( s11n::s11n_cast( ivec, dlist ) );
```

If the assertion succeeds, dlist contains a list of pointers to SomeT, copied from the objects in ivec. They could be int, char, MyType or whatever - any Serializable will do.

A generic implementation of s11n_cast() can be achieved in these few operations:

1. Create a temporary node.
2. Serialize the source Serializable into the temp node. On error return false.
3. Deserialize the node into the destination Deserializable and return result.

The actual implementation looks like:

```
template <typename NodeType, typename Type1, typename Type2>
bool s11n_cast( const Type1 & t1, Type2 & t2 ) {
    NodeType n;
    return serialize<NodeType,Type1>( n, t1 )
        && deserialize<NodeType,Type2>( n, t2 );
}
```

Again, i'm not saying this is a particularly *efficient* way to convert objects, but it is extremely generic. In theory it will work with *any* two types which use the same (or compatible) de/serialization algorithms. Out of the box, that's already millions of combinations, only counting STL-standard containers and PODs (that said, many non-STL containers work flawlessly with the STL-intended algos, as long as they follow the general published conventions).

28 Comparing s11n and Boost::serialization

This section tries to give an overview of the major similarities and differences between s11n and the only other serialization framework for C++ which can provide the range of the features s11n does: Dr. Robert Ramey's Boost serialization library, a member library of the Boost.org project. Below we will specifically address points and features which appear in either of s11n or Boost, but probably not in other libraries. Though "Boost" really refers to both an organization and the software that organization releases, here we will use the term Boost specifically to mean Robert's serialization library, which is part of the main Boost distribution as of version 1.3something (summer of 2004, if i recall correctly).

As a software library *user*, if i didn't have s11n, Robert's library would *definitely* be my choice for serialization support. If you are undecided on serialization libraries take a look at the Boost project, which provides not only

⁴⁷ What would the *Evangelism* section be without an *Amen* now and again?

serialization, but a huge number of industrial-strength libraries: <http://www.boost.org>

Please keep in mind that this chapter is *not* an attempt to sway you away from using Boost! On a coder level, i fully respect Robert's implementation and the design decisions he has made, and am *not* attempting to show that either library is significantly all-around better than the other. However, s11n has only one "competing" product, as far as i'm concerned, and i thought it might be interesting to compare them here. We will assume that the user is familiar with both s11n and Boost, or at least familiar with some of the main design aspects from both.

To open the comparisons on a positive note: Robert and i appear to agree on a great many design decisions. As his docs currently say about this library:

"Its has lots of differences - and lots in common with this implementation."

A quick comparison of the APIs would suggest that the projects two even co-developed at some point, though this is not the case⁴⁸.

28.1 Cans and cannots

Let's take turns listing a few features one lib has and the other does not, considering only out-of-the-box features which clients can get to by following the respective library manuals:

- Boost supports serialization of reference members in serializable classes, at least partially (the support might be fuller than the examples suggest). s11n does not directly support this.
- s11n supports loading without knowing the input format. Boost requires knowing the stream format and using the appropriate handler type.
- Boost internally tracks serialization of pointers and therefor inherently supports serialization of graphs. s11n requires client-written proxies to do this.
- Using Boost in client code effectively requires a hard dependency on much of the other Boost library, whereas s11n (as of 1.1) has no 3rd-party dependency requirements. Likewise, the boost.org libs provides a whole framework, whereas s11n provides only a serialization layer. (We will not count the STL as a dependency in either case because an STL implementation is required by most modern C++ code.)
- Boost directly supports serializing C-style arrays. s11n's author despises arrays and avoids them like the plague, but the framework theoretically supports them: use either a `for()` loop or a `for_each()` functor. The *nature* of both libraries' support is very different because of the fundamentally different pointer serialization policies.
- Boost provides several desirable features which s11n does not: `std::locale` and wide char/string support, `shared_ptr` support, and deserialization of classes containing references, to name a few.
- Likewise, s11n has a few interesting features which Boost does not: it overcomes some of Boost's current DLL-related limitations, supports transparent file de/compression, has more data handlers (3 formats in Boost vs. s11n's 8), and can do in-memory serialization (that is, without requiring intermediary i/o).

Most of these are relatively small differences or express clearly different design philosophies or even simply show a focus in a particular design *direction*. The overall range of features in both libraries is more or less comparable. i believe that both libraries can be used to implement most, if not almost all, features of the other with some relatively minor internal changes and the appropriate API wrappers.

28.2 Compiler and platform portability

Boost has s11n beat hands-down here. Robert has the major advantages of:

- A *lot* more experience than i with multiple platforms. My only development platform is Linux, with occasional access to a Solaris machine. In any case, my practical experience is limited to the GNU compiler and build tools. (That said, s11n is rigorously restricted to ISO-only C++ features.)
- The *massive* peer review effort which Boost.org is so famous for. This should *never* be underestimated.
- His software is built on top of other high-quality Boost software (e.g. Spirit does the file parsing), instead off of hand-rolled support code (e.g. the s11n file handlers are mostly implemented in flex-based parsers, rather unfortunately).
- One of Boost.org's core goals is platform-portable libraries. While i always try to adhere to published standards, and never use platform-specific constructs (except, of course, for platform-specific operations, like opening DLLs), i cannot personally test or support even a fraction of the platforms out

48 Robert, you interested? :)

there.

If your software already uses Boost, you should *strongly* consider using the Boost serialization library instead of s11n. i cannot confidently say that Boost-using code would benefit enough from s11n to justify the additional integration costs, considering that a good alternative solution is already available in Boost. While i do believe that s11n provides more features than Boost out of the box, i also believe that Boost could be made to do most, or even all, of the things s11n does with relatively little work. (i suspect that is a side-effect of their STL-ish architectures.) Even more specifically, i think that with the appropriate wrappers, the s11n and boost APIs could probably be made to effectively mimic one-another, at least where their features allow it, as their models are conceptually very similar and inherently very adaptable to this level of modification.

28.3 Archives vs S11nNodes

Boost uses an abstract "Archive" data store concept, which is fundamentally similar to s11n's S11n Node model. The main difference is that s11n separates the Node and i/o formats, where the Archive is a combination of data node and i/o marshaler. From a client level there would appear to be little difference in most cases. s11nlite explicitly abstracts away s11n's node type and i/o format, but i believe a similar wrapper would be trivial to add around the Boost code. Then again, the Boost API is simple enough that a wrapper like s11nlite is not really necessary.

Boost's approach is very similar to the model used by s11n's predecessor, which simply had a set of free functions for saving to or loading from the three different formats we had at the time. While it is straightforward and suitable for many purposes, i fundamentally feel that the only s11n-internal entity which should have to know about a stream's format is the code which reads and writes that specific grammar. Even the user shouldn't have to know what format he's using (admittedly, this is a purely philosophical standpoint, not a scientifically-backed one). Actually, the Archive type does not publish any stream-related APIs, even though they work similarly to streams. This means that they can be implemented to be grammar-neutral by simply adding another layer of indirection behind the existing Archiver interface or implementing your own Archiver which uses, e.g., a database as a back end.

s11n internally uses a factory interface for loading all i/o handlers, regardless of whether they are statically linked in with an application or are truly dynamically loaded via DLLs⁴⁹, and encourages users to not give a hoot about what data format they are actually using.

One perhaps-not-immediately-obvious advantage of s11n's approach is that it inherently provides the static approach as well as dynamic loading. That is, if you would like to specify a specific grammar handler there is nothing stopping you from doing so:

```
MyClass myobj;
...
s11nlite::node_type dest;
s11nlite::serialize( dest, myobj );
s11n::io::funxml_serializer ser;
ser.serialize( dest, std::cout );
```

And the converse for loading. You will need to include the proper Serializer header(s), of course, and possibly link against the Serializer (depends on how you build your libs11n). The more generic approach, and one which does not require the headers for each serializer is:

```
std::auto_ptr< s11nlite::serializer_interface >
    ser( s11nlite::create_serializer( 'funxml_serializer' ) );
if( ! ser.get() ){ ... damn ... }
ser->serialize( dest, std::cout );
```

While Boost does not currently appear to offer such a feature, i believe this is largely because the overall Boost project currently lacks a cohesive factory API, and this support could probably be added to Boost with relatively little work.

28.4 Non-intrusivity

Though our approaches are quite different, both libs provide functionally similar non-intrusive (i.e., proxied) serialization support. Robert's approach (via overloaded functions templated on the Archive type) is certainly more portable to older compilers than s11n's approach (mainly via template specializations). i must admit that i simply never thought of his approach before seeing his code, as s11n's model fit so well with template

⁴⁹ It is technically possible to write a classloader which literally creates the classes as needed, but i have never seen this implemented in C++ (the class creation/compilation overhead would be extreme, i think). It's been demonstrated in PHP, for example: creating database classes on-demand by analysing db table structures, creating class code to mimic them, and eval'ing it.

specializations that function overloads were simply never considered. In theory they can be used in conjunction with s11n's model, and *vice versa*. i cannot currently think of any reason why either approach would be fundamentally more or less powerful than the other, nor do they appear to be mutually exclusive in any way. Function overloads are certainly conceptually simpler, and probably *much* easier for new users to grasp, particularly those who are not well-versed in C++ templates.

28.5 Serialization of pointers

This is one of the points where, again, i admittedly stray far from conventional wisdom. Boost takes a very correct approach and has built-in support for tracking the addresses of serialized pointers, such that each is only serialized once and a graph can be correctly deserialized by the core library without user intervention or special support. Boost also has special support for `boost::shared_ptr<T>`, since that is a core component of the overall boost.org framework.

s11n differs quite radically, taking the "convenient" approach of simply treating serialized pointers as non-pointers. That is, serializing (T) and (T*) are functionally identically. During deserialization we rely on C++'s strong typing support to put us into a context where we can determine whether we need to deserialize a heap- or stack-based object. For example, deserializing data into a `list<T*>` will create T objects on the heap, whereas deserializing a `list<T>` will not. This type of difference is handled transparently by the library. The major cost for this is that it (probably) cannot provide built-in pointer tracking support for doing things like de/serializing graphs.

The separation of the core serialization API and i/o API in s11n make this even more difficult, as we need a data-format-agnostic way of building inter-node pointers, so to say. Again, this is a decision which i feel lies way outside of s11n's scope. For example, i don't want someone who uses s11n-generated XML in a non-s11n application to have to conform to the s11n-imposed conventions for embedding references to other nodes in the XML tree. Why not use a standard like those emerging from the W3C? Because s11n is data format agnostic and therefore doesn't know about *any* grammar standards. See the problem? i refuse to enforce force such a requirement on the base Serializer interface, as i feel it would greatly complicate their implementations. Having to write i/o parsers is bad enough as it is, and having to put that much more work into them doesn't sound like my idea of a fun coding session.

Serialization of graphs and other pointer-related tricks *can be and have been* done in s11n, but the core library provides no special support for them. Quite the opposite, the core goes out of its way to hide the differences of pointers and non-pointers!

28.6 Data Versioning

One fundamental design decision which needed to be made very early on in s11n's development was the issue of how to track versions of data layouts, such that we can tell if we are loading data with a different logical version and abort deserialization if we do.

This is another one of those points where i seem to disagree with every respectable programmer in the world. *Strongly* disagree, even. My decision was, and probably always will be:

*Data versioning support does not belong in this library's core. **Period.***

Of course, it's not fair to make such a strong blanket statement like that without backing up my case. Before i do, a short disclaimer is in order:

Libraries which do *not* use a key/value pair model for serializing class data *really do require* a built-in versioning system, and a lack of such support in these libraries would *indeed* be a problem. They write X data members to a stream and expect to be able to read X items from the stream, and need some core-accessible way of providing at least basic verification of that. Fair enough.

For reference purposes, let's call Boost's overall i/o approach the "X/X" (or "positional data") model, as it is inherently limited to the physical ordering of the serialized items. We could also call it the Ordered model, but "order" also has other implications which may or may not apply here. In any case, what distinguishes it from s11n, for our purposes, is that X/X requires data versioning to be built in to the core serialization library, whereas a key-value-pair (KVP) model does not.

My case against including this support in the s11n core boils down to the following:

- Doing so requires imposing "some sort of versioning conventions" on all clients. e.g. use incremental numbers or conventional software version numbers, like 1.2.3. This would have been an arbitrary design decision which s11n's author would have to impose on clients. The fewer such conventions the library imposes, the better.
- Doing so requires s11n to have some idea of what constitutes an incompatible version, potentially including support for version number comparisons to allow operations like "support up to 2 revs back"

- or "compatibility == the same major and minor numbers, irrespective of patch level" or other such oddness.
- How do we report versioning errors? Using the normal return-false approach or a special approach (e.g. version-related exceptions)? Again: that would be an arbitrary decision which s11n would impose on you.
 - My personal experiences has shown versioning to be a significant *hinderance*. This is probably because i code almost exclusively on Open Source projects, which inherently tend to fluctuate a lot more than commercial products do. (Mine do, anyway. ;)
 - The KVP model, e.g. as used by XML-based applications, appears to be far more version-flexible than people give it credit for. Data versioning can be implemented within this model at theoretically *any level* of a data tree - from the lowest integer member to the root-most node of a data tree, and it can be done independently of any data format. There are many different ways to implement this, both intrusive and non-intrusive, and it would not be fair for s11n to impose any specific implementation on you.
 - Never in my coding life (let's call that 10+ years, if it makes a difference) have i ever needed data versioning for proper function of my applications. If the user feeds us properly structured data, deserialization works, otherwise it fails. Why make it more complex than that? As in XML-based applications, semantic validation is necessarily a client-side choice and versioning falls into the category of semantic validation. s11n concerns itself with the *structure* of the data, and cares very little for the semantics of the data (and then only for classloading, because we have to store a unique-per-classloader identifier for each C++ class).
 - And finally...
Computers are inherently stupid, and the thought of a piece of *software* telling *ME* what data i am *permitted* to feed *MY* application makes me queasy. It makes me downright mad, actually. This is our decision to make, not s11n's, and s11n's architecture allows us to make such determinations at almost any given point in the deserialization process, should we want to.

A quick, incomplete comparison of the properties of each model reveals the following notable practical differences:

- The X/X approach is grammatically more compact, potentially *drastically* so. For proof of this just compare any XML file to the equivalent in a binary grammar. The addition of client-transparent stream compressors (e.g. built on top of zlib or bz2lib) makes this point largely moot, at least for practical purposes (though not techno-philosophically, because such features are not always readily available in all projects).
- The KVP approach writes *named* elements and can search for them later by name. Thus we can add properties, remove them, check under different names for the same property, and other operations related to version interoperability. That capability is not quite missing the X/X approach, because we can potentially map version numbers to specific deserialization operations, but we don't have the playroom which KVP allows for.
- The X/X approach would appear to require more maintenance than KVP-based code when a class gets new members. Robert's X/X implementation is quite sane, but still requires some amount of care on our part if we want to support older data files as our objects change, if only because (a) each developer has his own philosophies about version numbers and (b) the version number is defined at one source code point and accounted for at another point, which makes them easy to get out of sync, especially in multi-developer projects. In X/X, a failure to change a class' version number when its serialization algo changes (e.g. as data members are added or removed) can result in unpredictable, or even undefined, runtime behaviour. (i believe Boost explicitly throws if it detects this problem, but i am not 100% certain of that. RTFM.)
- The X/X approach possibly provides easier trackability of pointers when doing things like serializing graphs. Theoretically, though i can't really back that up at the moment. s11n's "deep pointer copy" policy shifts such "special-case" work to the client, whereas Robert's code handles all of this transparently.
- Data files created for X/X models are inherently unusable by KVP models, but the other way around is not the case because we can always discard name info later to create X/X data from a KVP data set. It is interesting to note that Robert's documentation shows an example of serializing using a KVP interface, in which the key is internally discarded.

Which approach is better, KVP or X/X? As always, it really depends on what your needs are. i obviously prefer the KVP approach, and personally consider details like data compactness to be "issues of the past" (so sue me - i almost always choose convenience over drive space).

28.7 API ease of use

Boost is probably much simpler to get started with than s11n is. Boost's public API is very straightforward, even almost intuitive. While s11n's public API is just as simple, s11n sets out to specifically abstract away a couple more details than Boost does and has a proportionally (perhaps even disproportionately) higher learning

curve. For example, Boost does not appear to have a public factory/classloading layer, so those details never come into play.

Once the learning curve is climbed, s11n and Boost have approximately the same ease-of-use, i think.

Boost also takes advantage of operator overloading to provide a simplified client-side API. For example, if A is some Archive object and S is some serializable object, you can probably guess what the following operations do:

```
A << S;
A >> S;
```

Fundamentally, this shouldn't be a problem to add to s11n. Practically, however, s11n's use of the `node_traits<>` type as an API marshaler for arbitrary node types complicates the matter, as the operators would really need to be part of that `node_traits<>` interface. While i haven't tried it out, i do not believe it would add to s11n's ease of use the same way it does in Boost, mainly due to having to create a traits object (or some middle-man) to apply the operators to. Constness of nodes complicates this - we would need two such types, one for const nodes and one for non-const, in order to hold a const-correct pointer/reference to the node. Tried that, and it was ugly. Also, s11n's strongly DOM-oriented data model complicates operations like the hypothetical:

```
typedef node_ioops<s11n::node_type> NIO;
NIO n( my_node );
n << object1; // we can't know the object's/node's name
n << object2; // same problem, plus...
n >> object3; // which object do we read?
```

Additionally, s11n's whole i/o model would inherently complicate such an addition, as discussed in section 23.5.4.

If a user is willing to stick with a single concrete data node type, such operators could of course be part of that API. i am not keen on the idea of adding them to the core node interface, however, even though in Boost's case i do consider them to be well justified.

28.8 Serialization Traits

That s11n and Boost both use traits types to store information about serializable types is pure coincidence. We both use them for tying metadata to types for purposes of managing serialization, but we do *completely* different things with them. Boost manages, for example, pointer tracking, custom RTTI [Run-Time Type Information], and data version number (a very clever place to put it, actually), whereas s11n mainly uses it for providing typedefs and (as of 1.1) access to class names (which is conceptually similar what Robert does with his RTTI).

It was by reading the Boost documentation that i learned that s11n's proxying and traits approaches will only *properly* work on C++ platforms which *fully/properly* support partial template specialization. On others it might not choose the proper specialized types. i have *no idea* what compilers might be troublesome here. Not mine, anyway ;). Again, this is a design choice of s11n: it requires a more modern compiler than Robert's library does.

28.9 Efficiency

Again, Boost has s11n beat hands down on this, on all accounts.

One of the reasons is that Boost uses parsers written using Boost::Spirit, a true wonder of technology which obsoletes tools like lex for C++ projects and generates code which compilers can theoretically optimize down to the last bit. The unfortunate fact is that most of s11n's input handlers are written in lex, and this includes a rather large amount of underlying support code to help lex code fit into the modern C++ world more satisfactorily. Except from the fact that it works quite well, the amount of lex support code is not something i'm proud of.

i would love to use Spirit, because it's just too cool to overlook: <http://spirit.sourceforge.net>. Initial tests with it revealed that Spirir (apparently) requires that all input be buffered. Experience has shown that such buffering is not a viable option as default Serializer behaviour because data is arbitrarily large (one user reported using a 1GB XML file).

To be clear, neither Boost nor s11n inherently rely on either Spirit or lex, or any other parser framework for that matter, but a serialization library without *some* form of included i/o support is pretty useless for most cases (but not all cases⁵⁰). This i/o support takes the form of some type of parser, but this is largely an implementation

⁵⁰ There are actually valid uses for serialization without any underlying i/o, like databases, shared-memory (where objects could be written directly), or even passing objects around via a clipboard-like mechanism.

detail and normally need not intrude on clients at all.

Another area where Boost is inherently much faster than s11n is in its one-pass de/serialization model. The Archive type *is* the i/o marshaler, and all de/serialization operations are performed directly on Archive objects. In s11n we de/serialize objects from/to containers, similar to how we would in an Archive, and it is these containers of "raw" data which are used by the i/o handlers. This is an unfortunate cost of the physical separation of core serialization operations and stream i/o, but one which i believe is highly justified for this library.

That said, it is theoretically possible to add internal i/o support to a new Data Node type and use that node type with s11n to provide similar functionality as Boost's Archive type. Likewise, it is theoretically possible to similarly wrap up Boost's Archive type to use two-phase de/serialization (as if you'd want to). Both architectures are very flexible to this type of change.

28.10 The interesting part is...

In hindsight (after having written this chapter, which included reading much of Robert's documentation and some of his source code), the following points have become clear to me:

- Robert and i are indeed, as he once said in an email, "kindred souls," both out just trying to save our objects.
- On the surface, s11n and Robert's code have a few similarities. All coincidental.
- At the overall architecture level, they have an *uncanny* number of similarities. Again, all coincidental.
- The implementation details are *completely* different animals.

That last point, in particular, strikes me because what's *really interesting* about it is: they are different animals *for completely different reasons*. That is, the features Robert's code and s11n provide are not necessarily mutually exclusive, but often exist either as different approaches to the same end or as solutions to completely different parts of the overall serialization process. In some cases each goes into areas the other simply has not explored. A couple examples include:

- Robert's Archive and s11n's data node models are not only both there to serve the same end (the client's interface to and from The Void), but also are both important templated types for the architectures.
- Robert's Serialization traits types track pointers, version numbers, and RTTI info, amongst other things. s11n's traits provide the `class_name()` function, which logically overlaps somewhat with the RTTI features, and several functors which play a similar role as function overloading does in Boost's code (and it does so in a mutually compatible way, it turns out).

The main implication of this would seem to be that it might be completely worthwhile to look at either merging in features from each other's library or to work out some way to merge them. A simple disappearance of one of the libs would not be acceptable by either of us, i'm certain, and i do feel that both distinguish themselves enough that they cannot simply merge one-to-one. It would be interesting to figure out how the core differences of, e.g. versioning and deep versus shallow pointer copying, could be abstracted into policies types or other C++ techniques, such that we could present a single core and build our own features on top of it. After having read much of Robert's documentation, i have little reason to think that this is not possible. The difficult part, i think, is figuring out where the line between core and client-side policies should come in. Something to think about, anyway...

28.11 In closing: s11n.net and Boost.org

To be clear, no s11n.net software has *any association whatsoever* with Boost.org's software, and we won't defame them by claiming any such association.

From here on we switch from "Boost" meaning "Robert Ramey's Boost serialization library" to Boost meaning the Boost.org libraries in general.

Several people have written me to ask if i plan on submitting s11n to Boost.org for consideration as a member library.

i'm *truly flattered* by this question, but i have no plans on submitting s11n to Boost.org. The reasons are:

(Please accept my appologies in advance if any of the reasons below seem presumtuous, pompous or even downright stupid. Everyone's got their own quirks, and a several of mine are expressed below.)

- They are Gods, i am not⁵¹. They would eat me alive and call it a Virgin Sacrifice Breakfast. i am a long-

⁵¹ "Ray, the next time somebody asks you if you're a *god*, say YES!" - Ghostbusters

time hobby programmer who's hobby serendipitously turned into not only his profession but also his lifestyle. A virtual hippie, so to say. By comparison, many of the Boost members are well-trained, seasoned veterans of far more design committees and software wars than i.

- i believe the Boost team would (quite rightfully) try to enforce a strict set of exception conventions on the library. As discussed in section 4.6 of this manual, i currently have reservations against doing so. i don't want to be faced with that reality quite yet. The day will likely eventually come, but only after i feel comfortable with all of the design decisions and their implications.
- i would likely be required to explicitly support, or help to support, a wide variety of C++ platforms which i will never in my life lay fingers on. i could not, in good conscience, possibly claim to support platforms who's names i know only from `#defines` in `config.hpp`⁵². Likewise, i despise spending a significant amount of my coding time researching workarounds for deficient platforms, even if i *do* have access to those platforms. i *love* software development, and i want it to continue being *play-time* instead of turning into *work-time*.
- My freedom to experiment in the main source tree would be more limited, as stabler interfaces would be that much more important. Either that, or i would end up maintaining two different copies. That would not only be a real drag, but would also send the wrong message to users by providing two potentially incompatible APIs.
- While there are some compelling differences between s11n and Robert's implementation, our libraries are uncannily similar in both nature and design. i believe that in bending s11n more towards Boost we would end up at roughly the same implementation, or at least very similar features wrapped up in very similar interfaces. Neither Boost.org nor its users would benefit from an overlap of that size, even if "some competition within Boost might be a good thing" (as one writer suggested). Boost is a cohesive whole, and non-duplication of features helps keep it that way.
- We could probably never get a group consensus agreeing to keep s11n's deep-pointer-copy policy (more likely, i would be outvoted 400 to 1). Nor would we ever find a 100% all-around-agreeable factory interface, including the underlying conventions. Nor would most efficiency-seekers even look twice at s11n's heavy use of lexical casting, would demand internal native type support via, e.g. `boost::Any`, would require strict performance definitions, etc. Fair enough, but that simply isn't my thing.

By and large, i'm worried about Death by Committee even more than the death by Virgin Sacrifice Breakfast, though i'm not sure who would die first, s11n or my desire to continue coding on it.

To be absolutely clear: both this library and i *would* certainly *both* benefit *greatly* from the Boost code review process⁵³! Well, the one of us who didn't die first would, anyway ;). i want to save my objects *now*, and s11n does that *now*... and does so *without killing anyone*⁵⁴!).

It is possible, but i don't quite dare say "likely", that i will at some point fork off a copy of s11n which is based off of the core Boost libraries, targeted specifically at Boost-using client code. This primarily depends on the availability of Boost on client machines (traditionally it is not preinstalled on most systems).

One of s11n's long-standing design decisions has been to reduce 3rd-party library dependencies to a minimum. Thus i spent 2+ years writing utility code which already exists in libraries like Boost :/. If we were to replace all of s11n's "utility code" with Boost equivalents, we could probably cut the size of the tree by 1/2, not counting the i/o parts (that makes up the majority of s11n's code). And i could *finally* get rid of that damned string utility library which keeps hopping from source tree to source tree like a little virus.

Assuming even a modest 20% code reduction, that would equate to 20% less code to maintain, which is always a good thing. Of course, it also means relying on gawd-only-knows-how-many underlying libraries in Boost, the interfaces and behaviours of which we can only hope are stable from one version to the next. (To be clear, i have no experience with Boost version compatibility, so i am not badmouthing them here!)

Not to be underestimated: some of the Boost code will theoretically become part the "next" C++ standard library and it would pay notable maintenance dividends to base s11n off of these libraries as much as possible.

i feel compelled to make a final confession, as well, and explain the reason why s11n is not already built off of the Boost libraries. This has been asked more than once, and the question is a fair one.

i have some deeply-seated, admittedly somewhat eccentric, philosophical problems with the Boost distribution policies. Not their licence, but the *way* their code is distributed.

In short, my message to the Boost team is this:

52 i'll save the *Tirade on the Illusions of Portability as Perceived by Most Autotools Users* for another time.

53 TODO: see if there's a Boost-supported process to submit code for review with the explicit idea that it is not targeted at inclusion for Boost. i suspect not, given the necessary overhead, but it would indeed be very interesting. A "Boost of Breed" stamp of approval type of thing.

54 If this *does* happen to you, please file a bug report.

If the code was easy to install, i would have been using Boost since years. Please provide some form of conventional build process (one that doesn't force me to download the build tools!). Whether or not they are Autotools, i don't care: a simple configure script and/or Makefile would do. Justification: as a library coder, if i do not believe that Library ABC will be on my target client systems, i generally will not introduce a dependency on Library ABC in my libraries. i'm pedantic about that, to the point of even skipping over jewels like Boost if their value isn't relatively convenient to cash in on.

And get rid of the `config.hpp` "feature" of `#erroring` on the unknown compiler version every time i upgrade my gcc!!!! ARGH!!!!

i admittedly get overly-annoyed when it comes to points like these, but if you guys will fix these things then i'm your newest convert for life. The wonderful code - and even complete documentation - is all there. Practically a C++ Nirvana right before our drooling mouths, but it is nonetheless not as accessible as it should be.

Potential Boost users: please pay no attention whatsoever to this man's ramblings - give Boost a try and you will probably be amazed by its quality and range of features.

29 Source tree innards

This section contain information about some of the implementation details of s11n, and is only of potential interest to those working directly with the s11n sources. It may be of particular interest to anyone attempting to port the tree to another platform.

29.1 Build tree structure

The build tree is structured in a fairly straightforward, mostly conventional manner. It looks more or less like this:

`toc/` = the complete build tools (toc means= "the other configure").

`doc/` = the docs (this file), plus possibly some Doxygen stuff.

`include/` = empty (just one Makefile). The headers get symlinked here during the build process.

`src/` = the source code, of course, made up off the following trees (listed in build/dependencies order):

`plugin/` = the plugins sub-lib. Note that it comes *before* s11n in the dependency chain (but this may change someday).

`s11n/` = the core library, including the classloader/factory API.

`io/` = core i/o code, several subdirectories (one for each specific Serializer class), and shared utility code for the Serializer build process (e.g., creating the lexers).

`lite/` = s11nlite and friends.

`client/` = client-side code.

`s11nconvert/` = utility to convert between any two Serializers' formats.

`sample/` = client-side demo/sample/test code.

The `src` directory is broken down the way it is mainly to enforce specific dependencies between certain parts of the framework. For example, the core should never know about the i/o layer, and is thus built before the i/o parts (before the i/o headers are in place), to enforce this dependency. If someone accidentally adds `#include <s11n.net/s11n/io/...>` to a source file under `src/s11n`, the next full build would fail to compile (unless per chance the compiler picks up an *installed* copy of the header from, e.g., the build's \$prefix path).

29.2 Header file weirdness

All header files are stored in the same directory as their source file (if any, otherwise the same directory as their sub-library), but they are always referenced in other files using the fully-qualified form: `#include <s11n.net/s11n/...>`. This works because the headers are symlinked into place (under `include/s11n.net/...`) during the build process. This serves the following purposes:

- Enforces that we cannot build library A before library B if lib A depends on B, because B's headers will not be in place (and therefore, presumably A will not compile). This enforcement only works on initial/clean builds, by the way.
- Establishes the same `#include` convention that client code should use when including the project headers.
- It gives maintainers more flexibility and simplifies porting the tree to other platforms, as we can move the sources and headers around without breaking any `#includes`.
- While the source tree might contain "extra" headers, a coder can know which ones are "official" by looking in the include tree (which contains only headers which would get installed).

While this might seem odd, i've been using this approach since last millennium and it has always served me well.

We could just as well store the physical headers under `include/...`, but in my experience this makes editing the code more tedious. i prefer to have the headers and implementations in the same directory, and the symlinking provides that "extra layer of indirection" so that both approaches are accommodated simultaneously.

i've worked on several projects which split the sources and headers, and almost always find that coders inadvertently include headers from modules which come after their own in the dependencies chain. While this does not unduly upset most people, it does unduly upset me (i'm a huge fan of proper dependencies). There is no simple, straightforward way to find this type of problem in such a tree, so i prefer to make it impossible for a coder to do, via the symlink approach.

29.3 Generated files

The build tree includes the following generated files, which are normally created during the configure process. For porting purposes, they can be hand-created or taken from a system with a generated copy and tweaked to suit.

- `src/s11n/s11n_config.hpp`: this is the library's main configuration header, defining what features are supported, shared paths, etc. It is generated from the file `s11n_config.hpp.at`.
- `src/plugin/plugin_config.hpp`: this header is only needed on platforms where `s11n_config.hpp:S11N_CONFIG_ENABLE_PLUGINS` is set to a true value. It defines the plugin layer's options, such as the default plugins search path, and is generated from the file `plugin_config.hpp.at`.
- Various flex-generated lexers, `src/io/*/*.flex`. The tree ships with pre-flexed versions, however. In fact, the flex-generated versions won't even compile as-is under newer C++ compilers due to stricter C++ standards compliance in modern tools. The generated copies are hacked a bit during the build process using Perl, but this is only known to work for lexers generated by flex 2.5.4. While newer flex versions exist, Linux distributions ship with 2.5.4 since *years* because it at least generates compilable C code (whereas newer ones often fail to do even that).

29.4 Plugins

Platforms which meet the following requirements can potentially work with s11n's plugins model:

- Must have the equivalent of `dlopen()`. That is, a function (probably with a C API) which can open a DLL and link it into the running process.
- It must be able to export symbols in a DLL into the application. On GNU platforms this is done using the `-rdynamic` (or `-export-dynamic`) flag, and needs no special code support. On Windows platforms, all "appropriate" classes must be *explicitly* exported. This is a real bummer, and i cannot personally tell you which classes need it and which do not (because my platform doesn't need this). The file `s11n/export.hpp` defines the `S11N_EXPORT_API` macro, which is intended to be placed in the declaration for classes which need to be exported.
- When using Microsoft(tm) compilers (and maybe others), all DLLs built for this framework must be linked with the "keep unreferenced data" option. This is essential for factory registration to work. If this option is not used, the build will work but no factory registrations will happen - the end effect is that we cannot load new types via the factory API.

If your platform supports any of the following DLL loaders, the provided plugin implementations should be okay for use as-is on your system:

- `dlopen()` - the *de facto* Unix standard.
- `lt_dlopen()` - a GNU variant of `dlopen()`, ported to many platforms. This variant is chosen by the configure script if it is found, taking precedence over `dlopen()`.

- `LoadModule()` - the Win32 equivalent of `dlopen()`.

For supporting other loaders, see the file `src/plugin/plugin.cpp` for how the platform-dependent code is handled.

30 In Hindsight...

"Don't you look at me that way!"

Mom

"Hindsight is always 20/20."

Common proverb

This section is mainly a place for me to blab about specific elements of the library that i would like to change, see changed, or "would/should have done differently." This is not a bug list, but might partially be considered an RFE (Request For Enhancements).

30.1 The name "Data Node"

This was a huge mistake. When the templated Node concept entered the API, i already had a type named `s11n_node` (but not the same one we have today), and didn't want to use the concept name `S11nNode` because i didn't want to give the impression that `s11n_node` and `S11nNode` were the same thing. Let's chalk up one point for Laziness. In hindsight, i should have thought more about it and chosen a completely different name, like `SerializationNode` (`SNode`, for short). Ashamedly, *that* name never hit me until writing this.

The phrase "data node" is simply too vague, and often ambiguous (e.g., in the context of serializing a graph, where "node" is the conventional term for each graph element).

In the future i may well start to replace the term. The fact is, however, that this document, the API docs, and the web site, are all *filled* with the phrase "data node". The effort needed to completely update the docs would be tremendous. i have reservations about "slowly" switching terms, though, because i don't want the different terms to confuse users.

30.2 Patterns, formality, etc.

i think it's understandable that i never had any clue that this project would grow to the size it has. It started out life back in 2001 as a set of utility code which i knew i would need in order to eventually implement serialization. The library itself, as a formal entity, has evolved steadily, often rapidly, since late 2003. Unfortunately, i have always been so focused on playing with the code that i have neglected some formalities which would not only make users' lives easier, but would also help to improve the library. While i have been quite diligent about documentation, i haven't, until recently, begun to think of the library in terms of Patterns (see section 4.7). This is probably a side-effect of me being so buried in the implementation that i simply haven't stood back long enough to see the various Patterns. i hope to be able to document these more fully in the future, and perhaps even adjust some "non-Patterned" parts of the architecture where it seems that a particular Pattern would work well.

The authors of the book *C++ Template Metaprogramming* [CTM2005], David Abrahams and Aleksey Gurtovoy, claim that types like `s11n_traits<>`, which they describe as "blobs", are actually "anti-patterns", meaning "don't do that!" i feel that their position is well-justified within the context of their Metatemplate Programming Library (MPL) work, but not in the general case. The "blob" pattern does have its drawbacks, but also fills numerous roles very nicely.

30.3 Exceptions

As documented elsewhere in this manual, the exceptions support in versions prior to 1.1.3 was completely broken. To be fair, it wasn't designed to deal with exceptions until 1.1.0, and even then the handling code was far from adequate. The lack of strong exception guarantees was not a reflection of my ignorance of what exceptions *are*, but of my uncertainty about how to best to accommodate for them in C++. My preconceptions of exceptions stem from my Java years, but i am fully aware that exceptions in Java and C++ are different beasts, *and* fully aware that i don't know what all of those difference are. Knowing that there were lots of pitfalls to exception handling, i cautiously avoided the topic for some time. This is rapidly being remedied in the 1.1.3+ releases.

The exceptions support would have been done a lot earlier if i had not delayed implementing the `s11n::cleanup_serializable()` mechanism. The prototype for that was developed almost a full year before i included it in s11n. i was initially afraid that the additional overhead would add to the already-hurting client-side compile times. This fear turns out to have been unjustified - the impact is measurable but small. On one quick test it appeared to add about 1/3rd of a second to compile times per input file, though this number is actually dependent on the number of registered Serializable types. The benefit of that mechanism is immeasurable though, as it empowers many safety guarantees this library could not otherwise make. i personally don't mind paying 1/3rd of a second for the guaranty of no leak if an exception is thrown.

30.4 Build tree and code layout consistency

i know it's annoying that every 3rd release i move header files around. The fact is, i'm a habitual tinkerer. As i use the library in more client code, i change the library to be more accommodating, or just clearer or simpler to use.

This isn't likely to change.

Since the 1.0 release, the project officially has "stable" and "development" source trees, so i no longer feel guilty about this. Having the dev tree around keeps me from mucking up the stable interfaces, as i undoubtedly would if i didn't have a second branch of the source tree to freely experiment with.

31 Is this the end?

"How far y'all going, she asked with a sigh. We're goin' all the way. Til the wheels fall off and burn."

Bob Dylan, Brownsville Girl

We are nearing the end of the document, but hopefully the new possibilities for saving your data have just begun. :)

If you are looking for more information about using s11n, try:

- The s11n source tree has code for a couple client-side apps, which will certainly prove informative to those starting out with s11n: see `src/client/sample`
- The web site is updated fairly often, and you just might find something interesting over on there if you check back once in a while:
`http://s11n.net`
- If you have questions, concerns, or just want to say "Hello, world", please email us:
`s11n-devel@lists.sourceforge.net`

.....

Before i go, i want to tell you briefly why *i* use s11n in all of *my* code: because it's just so damned easy to do. When there are such time- and feature-gains to be had via such a simple-to-integrate tool, it's hard to justify re-implementing any save/load code⁵⁵. This continual interaction with multiple clients also greatly helps in figuring out exactly what s11n needs to do and what services it must provide, so the library continually reshapes and improves under the well-proven and very-very-very long-standing rules of Natural Selection, also known as Darwinistic Processes or, in the marketing department, Upgrades.

As always:

- The source tree is *a/ways* the most-definitive source of information, but the web site is also updated fairly often as new advances are made, often a bit in advance of upcoming changes.
- i am always open to getting mails with questions about s11n, so don't hesitate to email our development list. i will ask that you please browse the manual first, but i certainly do *not* expect you to scour every web page or code file before posing a question. i understand that the documentation has some gaping holes in it, and i will be happy to fill those holes by answering your questions.
- The main goal of s11n is to *Save Our Data!* If s11n can't do that, please help us out by suggesting how we might be able to change it so that it *can* save your data! Sometimes just saying "s11n can't do [this]" is enough to spur a solution, as often the author does not realize something is a problem or omission until someone else points it out (thanks again to Ton and Gary, especially, for that).

Once again: thanks *a lot* for taking the time to consider adding s11n to your toolkit! And thanks *a whole lot* for Reading The Full Manual. :)

⁵⁵ You can bet your emacs that i'm pretty sick of that part by now ;).

--- stephan@s11n.net

or, of course:

s5n@s11n.net

:)

Happy hacking!!!

32 Bibliography

CTM2005

C++ Template Metaprogramming, by David Abrahams and Aleksey Gurtovoy.

CCS2005

C++ Coding Standards, by Herb Sutter and Andrei Alexandrescu.

C++StandardLib

The C++ Standard Library (A Tutorial and Reference), by Nicolai Josuttis. Without a doubt the single most-used C++ book i own.

EffectiveC++2

Effective C++, 2nd Edition, by Scott Meyers.

EffectiveC++3

Effective C++, 3rd Edition, by Scott Meyers.

MoreEffectiveC++

More Effective C++, by Scott Meyers.

EffectiveSTL

Effective STL, by Scott Meyers.

Gotchas

C++ Gotchas, by Stephen C. Dewhurst.

33 Index

algorithms.....	
commonly used.....	54
definition.....	22
serialization.....	53
architecture, overview of.....	22
credits.....	7
proxies.....	
serialization.....	53
Serializable Types.....	
abstract types.....	103
Types.....	
Base Types.....	21, 50
creating a Serializable Type.....	44
Interface Types.....	21, 50
type traits.....	31
zlib.....	86

Below are old entries imported from the original document. These will/should be ported to OO-native entries.

Base Types, abstract

26.3

bool, as return type

t4.6

bool, justifying

4.6

- brute force deserialization
 - 10.5.3
- bz2lib
 - 22.7
- casting Serializables
 - 10.4.1
- caveats
 - 2.4 | 23.1
- class_name()
 - 15 | 15.1
- classloader, definition of
 - 4.1
- classloader, role in s11n
 - 4.2
- cloning Serializables
 - 22.5
- common problems
 - 26
- credits
 - 1.4
- cycles
 - 23.2
- Data Node, definition of
 - 4.1
- Data Node, setting class name
 - 5.3.1
- Data Nodes, class names of
 - 5.3
- Data Nodes, property key requirements
 - 4.4
- deserialization, brute force
 - 10.5.3
- deserialization, process
 - 4.3.2
- deserialize, definition of
 - 4.1
- deserializing objects
 - 10.5
- Disclaimers
 - 1.2
- elem_t (sample Serializable)
 - 11.2.1
- elem_t_s11n (sample proxy)
 - 11.2.1
- features, primary
 - 2.3
- feedback, providing
 - 1.3
- file extensions
 - 14.1.1
- formats, data
 - 14
- functor, definition
 - 4.1
- functors, serialization
 - 13
- graphs
 - 23.2
- impl_class()
 - 15 | 15.1
- indentation, Serializers and
 - 14.1.2
- Interface, Default Serializable
 - 4.1
- interfaces, cooperating with remote
 - 5.4

- interfaces, custom Serializable
 - 8.2
- License
 - 1.1
- magic cookies
 - 14.1.4
- Node Traits
 - 4.1
- node_traits
 - 4.1
- node_traits<>
 - 6.1
- nodes, finding children
 - 10.3
- ODR
 - 4.1
- One Definition Rule
 - 4.1
- operator, deserialize
 - 4.1 | 5.2 | 11.1.4
- operator, serialize
 - 4.1 | 5.1 | 11.1.3
- Patterns
 - 4.7 | 30.2
- problems, common
 - 26
- properties, error checking
 - 10.2.1
- properties, getting
 - 10.2
- properties, setting
 - 10.1
- proxies
 - 8.3 | 13
- proxies, commonly used
 - 13.2
- proxies, specifying functors
 - 8.3
- proxy, list_serializer_proxy
 - 13.1.2
- proxy, map_serializer_proxy
 - 13.1.4
- proxy, pair_serializer_proxy
 - 13.1.5
- proxy, streamable_type_serialization_proxy
 - 13.1.1
- proxy, value_map_serializer_proxy
 - 13.1.3
- registration, class names
 - 12.3
- registration, custom Serializable interfaces
 - 12.5
- registration, default interface
 - 12.4
- registration, proxies
 - 12.6
- registration, where to do it
 - 12.8
- s11n, meanings of
 - 4.1
- s11n_cast
 - 22.4
- s11n_cast()
 - 10.4.1
- S11N_DESERIALIZE_FUNCTOR
 - 9

- S11N_SERIALIZE_FUNCTOR
 - 9
- s11n_traits
 - 4.1
- s11n_traits<>
 - 6.2
- S11N_TYPE
 - 9
- S11N_TYPE_NAME
 - 9
- s11nconvert
 - 21.1
- s11nlite
 - 2.5
- s11nlite, role in s11n
 - 4.2
- SAM
 - 4.1 | 17
- SAM, overview
 - 4.2
- Serializable interface, conventions
 - 5
- Serializable Traits
 - 4.1
- Serializable type, creating
 - 8.1 | 11
- serializable, definition of
 - 4.1 | 4.1
- Serializables, abstract
 - 26.3
- Serializables, casting
 - 10.4.1
- Serializables, creating
 - 8
- Serializables, working with
 - 10
- Serialization API Marshaling
 - 17
- serialization operators, templates as
 - 5.5
- serialization, process
 - 4.3.1
- serialize, definition of
 - 4.1
- Serializer, compact
 - 14.2.1
- Serializer, definition of
 - 4.1
- Serializer, expatxml
 - 14.2.2
- Serializer, funtxt
 - 14.2.3
- Serializer, funxml
 - 14.2.4
- Serializer, parens
 - 14.2.5
- Serializer, simplexml
 - 14.2.7
- Serializers
 - 14
- Serializers, conventions
 - 14.1
- Serializers, in s11nlite
 - 14.3.2
- Serializers, role in s11n
 - 4.2

- serializing objects
 - 10.5
- serializing Streamable Types
 - 10.4
- state, saving application-wide
 - 22.2
- Streamable Types
 - 10.2.2
- Streamable Types, definition of
 - 4.1
- Streamable Types, serializing
 - 10.4
- Streamables
 - 10.2.2
- Style Points
 - 4.1
- Supermacros
 - 12
- terms and definitions
 - 4.1
- thread safety
 - 23.3
- Traits, Serializable
 - 4.1
- type traits
 - 6
- 11
- 22.7